ConneXions —
The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

From the Editor

Welcome to Paris and *NetWorld+Interop 94,* the fifth and final stop on our 1994 World Tour. This issue of *ConneXions* is being provided to all conference, tutorial and exhibition attendees while supplies last. We are pleased to offer a special 20% discount on all new subscriptions. To take advantage of this offer, simply complete and return the order form on the back. For those of you unfamiliar with this publication, *ConneXions* is a monthly technical journal covering all aspects of computer networking and interoperability. It is the companion journal to the NetWorld+Interop conference, and has been published by Interop Company since the first Interop event in March 1987. For a complete index of back issues, send e-mail to: **connexions@interop.com.**

Our first article is the second installment in a series on cryptographic algorithms for providing security in computer networks. William Stallings describes public-key cryptography and secure hash functions. (Part One appeared in our September issue.)

Asynchronous Transfer Mode (ATM), also known as Broadband ISDN (B-ISDN) is probably the most important technology to emerge in the last few years. At NetWorld+Interop 94 you will hear a lot about ATM, both in the conference/tutorial program and on the tradeshow floor. *ConneXions* continues to cover issues related to ATM. Our second article by Reto Beeler of Ascom Tech AG describes a way to support data services in a connection-oriented mode, in contrast to the more typical connectionless mode.

Videoconferencing used to be limited to special-purpose equipment and data circuits, but in recent years it has extended its reach to low-cost personal workstations and wide-area computer networks. In the Internet, several audio and video conferencing systems have been in use for some time. One widely used conferencing system is *IVS* developed by the Institut National de Recherche en Informatique et en Automatique (INRIA). The system is outlined here by Thierry Turletti.

Last month, Jim Galvin of Trusted Information Systems gave an overview of the various activities related to security within the *Internet Engineering Task Force* (IETF), the primary standards development body for the Internet. This month, Dr. Galvin gives us a quick update from the most recent IETF meeting.

You can find more information on all of these topics in our extensive conference and tutorial program, and we hope you will continue to receive updates on emerging technologies through a subscription to *ConneXions—The Interoperability Report.* Enjoy your week in Paris!

Back to Basics:
Cryptographic Algorithms
Part II: Public-key Encryption and Secure Hash Functions
by William Stallings

Introduction

A growing proportion of the applications and protocols used over the Internet either have significant security-related features or have as their primary purpose the provision of some security facility. At the application level, examples include e-mail security (Privacy Enhanced Mail, PEM; Pretty Good Privacy, PGP), network management (Simple Network Management Protocol version 2, SNMPv2), and remote authentication (*Kerberos*). A common feature of all of these applications and protocols is the use of cryptographic algorithms to implement particular security services. The many such algorithms in use fall into three categories: conventional encryption algorithms, public-key cryptography algorithms, and secure hash functions. Last month's article dealt with conventional encryption and provided an overview of important algorithms in each category. [7] This month, public-key cryptography and secure hash functions are covered.

Public-Key Encryption

Public-key encryption, first publicly proposed by Diffie and Hellman in 1976 [1], is the first truly revolutionary advance in encryption in literally thousands of years. For one thing, public-key algorithms are based on mathematical functions rather than on substitution and permutation. But more important, public-key cryptography is asymmetric, involving the use of two separate keys, in contrast to the symmetric conventional encryption, which uses only one key. The use of two keys has profound consequences in the areas of confidentiality, key distribution, and authentication.

Misconceptions

Before proceeding, we should mention several common misconceptions concerning public-key encryption. One such misconception is that public-key encryption is more secure from cryptanalysis than conventional encryption. In fact, the security of any encryption scheme depends on the length of the key and the computational work involved in breaking a cipher. There is nothing in principle about either conventional or public-key encryption that makes one superior to another from the point of view of resisting cryptanalysis. A second misconception is that public-key encryption is a general-purpose technique that has made conventional encryption obsolete. On the contrary, because of the computational overhead of current public-key encryption schemes, there seems no foreseeable likelihood that conventional encryption will be abandoned. Finally, there is a feeling that key distribution is trivial when using public-key encryption, compared to the rather cumbersome handshaking involved with key distribution centers for conventional encryption. In fact, some form of protocol is needed, often involving a central agent, and the procedures involved are no simpler nor any more efficient than those required for conventional encryption.

Characteristics

A public-key cryptographic algorithm relies on one key for encryption and a different but related key for decryption. Furthermore, these algorithms have the following important characteristic:

- It is computationally infeasible to determine the decryption key given only knowledge of the cryptographic algorithm and the encryption key.

In addition, some algorithms, such as RSA, also exhibit the following characteristic:

- Either of the two related keys can be used for encryption, with the other used for decryption.

The essential steps are the following:

1. Each end system in a network generates a pair of keys to be used for encryption and decryption of messages that it will receive.

2. Each system publishes its encryption key by placing it in a public register or file. This is the public key. The companion key is kept private.

3. If A wishes to send a message to B, it encrypts the message using B's public key.

4. When B receives the message, it decrypts it using B's private key. No other recipient can decrypt the message because only B knows B's private key.

With this approach, all participants have access to public keys, and private keys are generated locally by each participant and therefore need never be distributed. As long as a system controls its private key, its incoming communication is secure. At any time, a system can change its private key and publish the companion public key to replace its old public key.
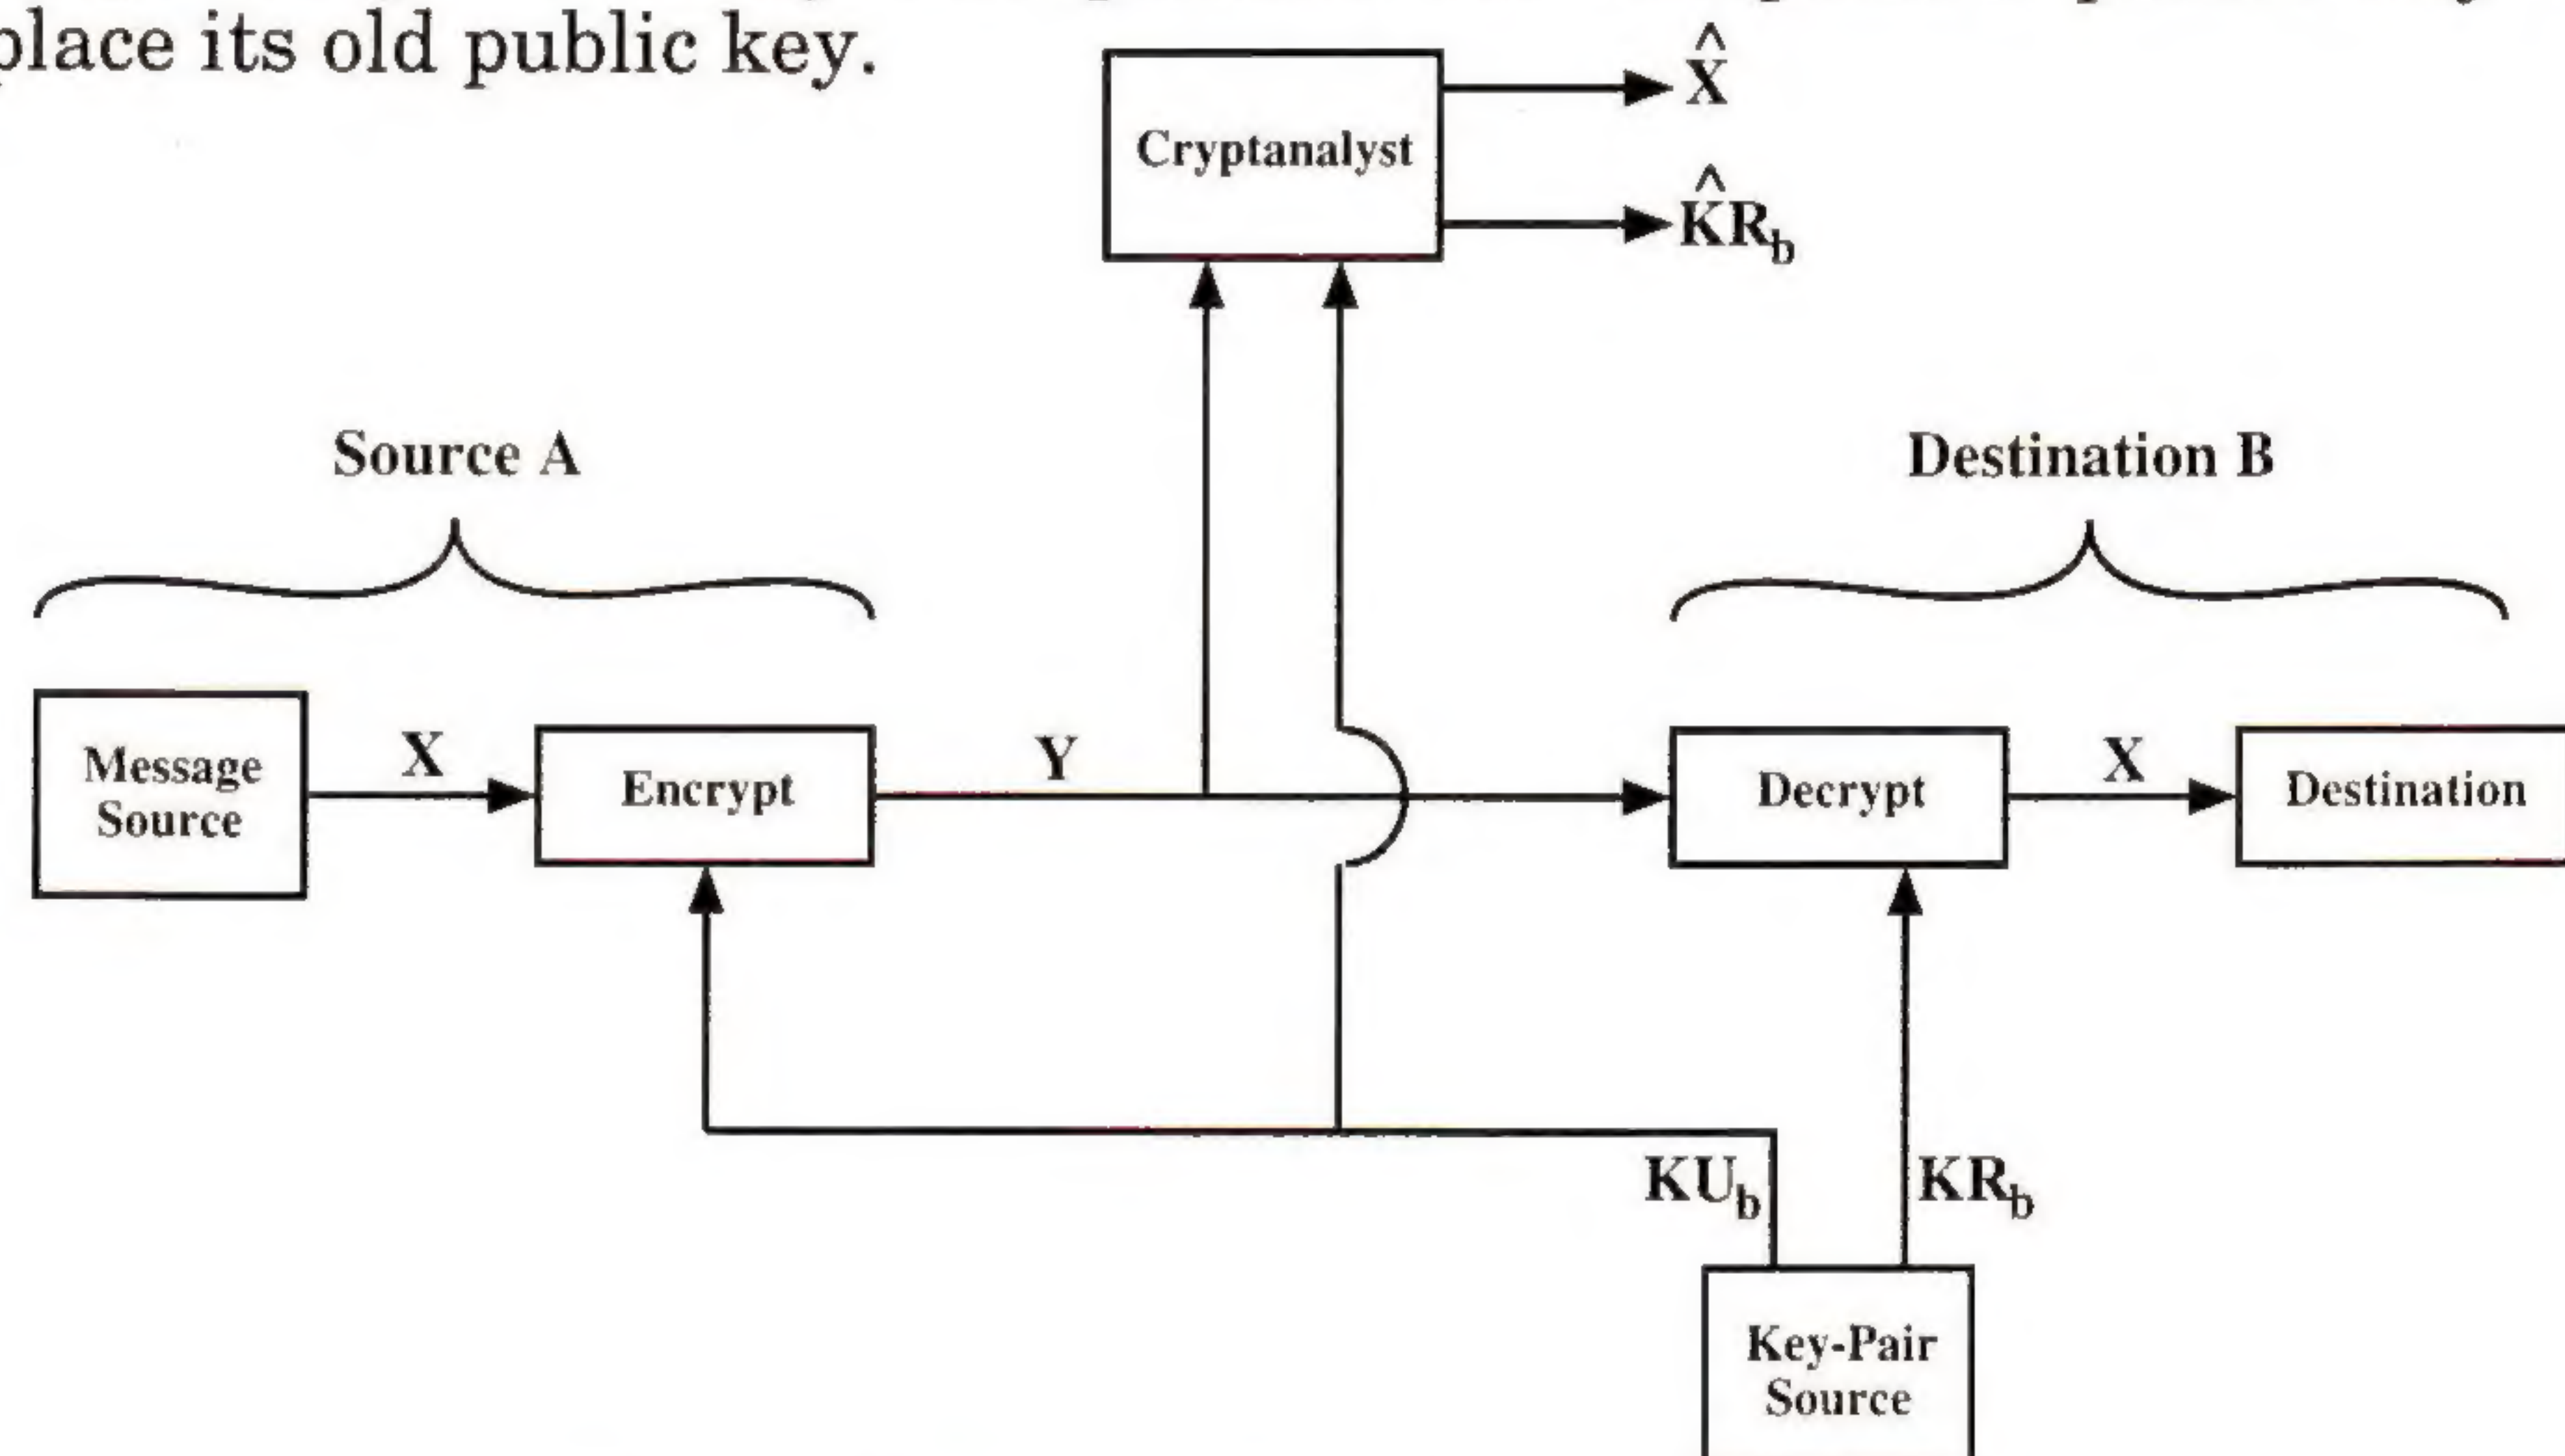

Figure 1: Public-Key Cryptosystem

Figure 1 illustrates the process. There is some source A for a message, which produces a message in plaintext, $X = [X_1, X_2, \ldots, X_M]$. The M elements of X are letters in some finite alphabet. The message is intended for destination B. B generates a related pair of keys: a public key KU_b, and a private key, KR_b. KR_b is secret, known only to B, whereas KU_b is publicly available, and therefore accessible by A.

With the message X and the encryption key KU_b as input, A forms the ciphertext $Y = [Y_1, Y_2, \ldots, Y_N]$:

$$Y = E_{KUb}(X)$$

The intended receiver, in possession of the matching private key, is able to invert the transformation:

$$X = D_{KRb}(Y)$$

An opponent, observing Y and having access to KU_b, but not having access to KR_b or X, must attempt to recover X and/or KR_b. It is assumed that the opponent does have knowledge of the encryption (E) and decryption (D) algorithms.

3

Cryptographic Algorithms *(continued)*

Digital Signature

We mentioned earlier that either of the two related keys can be used for encryption, with the other being used for decryption. This enables a rather different cryptographic scheme to be implemented:

$$Y = E_{KRa}(X)$$

$$X = D_{KUa}(Y)$$

In this case, A prepares a message to B and encrypts it using A's private key before transmitting it. B can decrypt the message using A's public key. Because the message was encrypted using A's private key, only A could have prepared the message. Therefore, the entire encrypted message serves as a *digital signature*. In addition, it is impossible to alter the message without access to A's private key, so the message is authenticated both in terms of source and in terms of data integrity.

In the preceding scheme, the entire message is encrypted, which although validating both author and contents, requires a great deal of storage. Each document must be kept in plaintext to be used for practical purposes. A copy also must be stored in ciphertext so that the origin and contents can be verified in case of a dispute. A more efficient way of achieving the same results is to encrypt a small block of bits that is a function of the document. Such a block, called an *authenticator,* must have the property that it is infeasible to change the document without changing the authenticator. If the authenticator is encrypted with the sender's private key, it serves as a signature that verifies origin, content, and sequencing. A secure hash code, discussed later, can serve this function.

It is important to emphasize that the encryption process just described does *not* provide confidentiality, i.e., the message being sent is safe from alteration but not safe from eavesdropping. This is obvious in the case of a signature based on a portion of the message, since the rest of the message is transmitted in the clear. But even in the case of complete encryption, there is no protection of confidentiality since any observer can decrypt the message by using the sender's public key.

RSA

One of the first public-key schemes was developed in 1977 by Ron Rivest, Adi Shamir, and Len Adleman at MIT, and first published in 1978 [2]. The *Rivest-Shamir-Adleman* (RSA) scheme has since that time reigned supreme as the only widely accepted and implemented approach to public-key encryption. RSA is a block cipher in which the plaintext and ciphertext are integers between 0 and $n - 1$ for some n. Encryption and decryption are of the following form, for some plaintext block M and ciphertext block C:

$$C = M^e \bmod n$$

$$M = C^d \bmod n = (M^e)^d \bmod n = M^{ed} \bmod n$$

Both sender and receiver must know the value of n. The sender knows the value of e, and only the receiver knows the value of d. Thus, this is a public-key encryption algorithm with a public key of KU = {e, n} and a private key of KR = {d, n}. For this algorithm to be satisfactory for public-key encryption, the following requirements must be met:

1. Is it possible to find values of e, d, n such that M^{ed} = M mod n for all M < n?

2. Is it relatively easy to calculate M^e and C^d for all values of M < n?

3. Is it infeasible to determine d given e and n?

The answer to all the first two questions is "yes." The answer to the third question is "yes," for large values of e and n.

Figure 2 summarizes the algorithm. Begin by selecting two prime numbers, p and q and calculating their product n, which is the modulus for encryption and decryption. Next, we need the quantity $\phi(n)$, which is referred to as the *Euler totient* of n, which is the number of positive integers less than n and relatively prime to n. Then select an integer d that is relatively prime to $\phi(n)$, i.e., the greatest common divisor of d and $\phi(n)$ is 1. Finally, calculate e as the multiplicative inverse of d, modulo $\phi(n)$. It can be shown that d and e have the desired properties.

Key Generation

Select p, q p and q both prime

Calculate $n = p \times q$

Calculate $\phi(n) = (p-1)(q-1)$

Select integer d $\gcd(\phi(n), d) = 1; \ 1 < d < \phi(n)$

Calculate e $e = d^{-1} \bmod \phi(n)$

Public key $KU = \{e, n\}$

Private key $KR = \{d, n\}$

Encryption

Plaintext: $M < n$

Ciphertext: $C = M^e \pmod n$

Decryption

Ciphertext: C

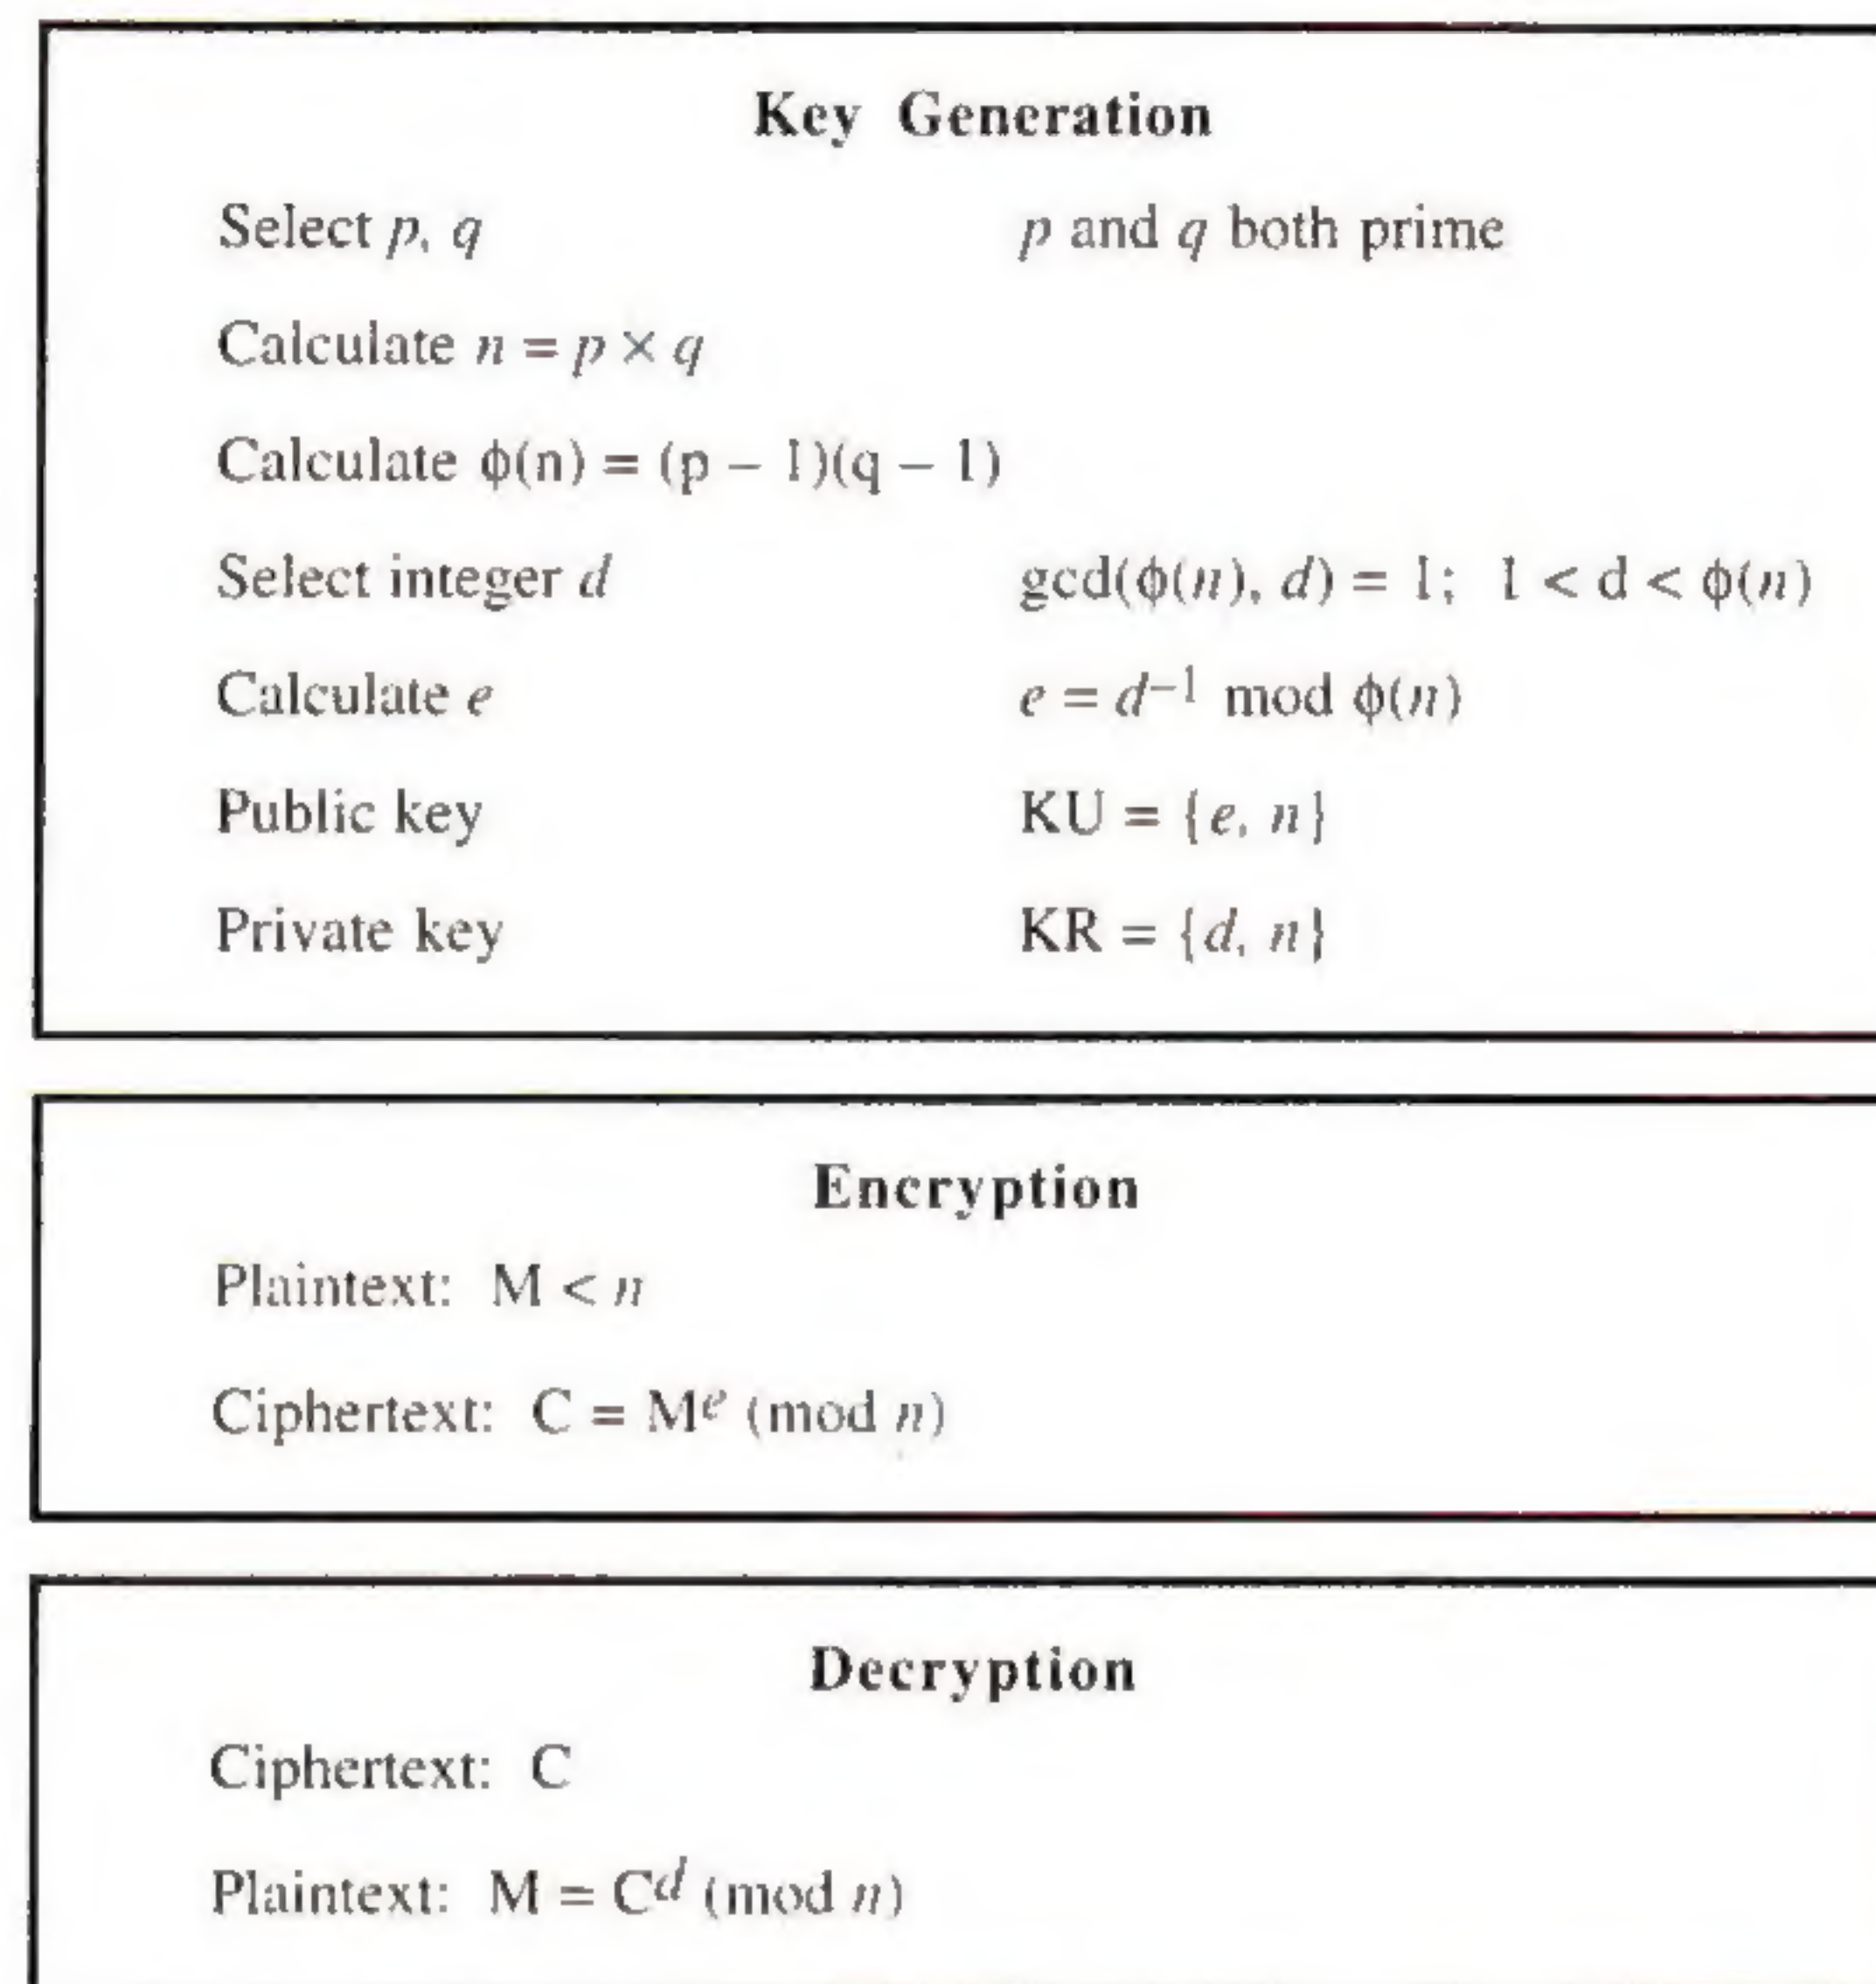Plaintext: $M = C^d \pmod n$

Figure 2: The RSA Algorithm

The private key consists of $\{d, n\}$ and the public key consists of $\{e, n\}$. Suppose that user A has published its public key and that user B wishes to send the message M to A. Then, B calculates $C = M^e \pmod n$ and transmits C. On receipt of this ciphertext, user A decrypts by calculating $M = C^d \pmod n$.

Example   An example is shown in Figure 3 on the next page. For this example, the keys were generated as follows:

1. Select two prime numbers, $p = 7$ and $q = 17$.

2. Calculate $n = pq = 7 \times 17 = 119$.

3. Calculate $\phi(n) = (p-1)(q-1) = 96$.

4. Select e such that e is relatively prime to $\phi(n) = 96$ and less than $\phi(n)$; in this case, $e = 5$.

5. Determine d such that $de = 1 \bmod 96$ and $d < 96$. The correct value is $d = 77$, because $77 \times 5 = 385 = 4 \times 96 + 1$.

The resulting keys are public key KU = {5, 119} and private key KR = {77, 119}. The example shows the use of these keys for a plaintext input of M = 19. For encryption, 19 is raised to the 5th power, yielding 2476099. Upon division by 119, the remainder is determined to be 66. Hence $19^5 \equiv 66 \bmod 119$, and the ciphertext is 66. For decryption, it is determined that $66^{77} \equiv 19 \bmod 119$.
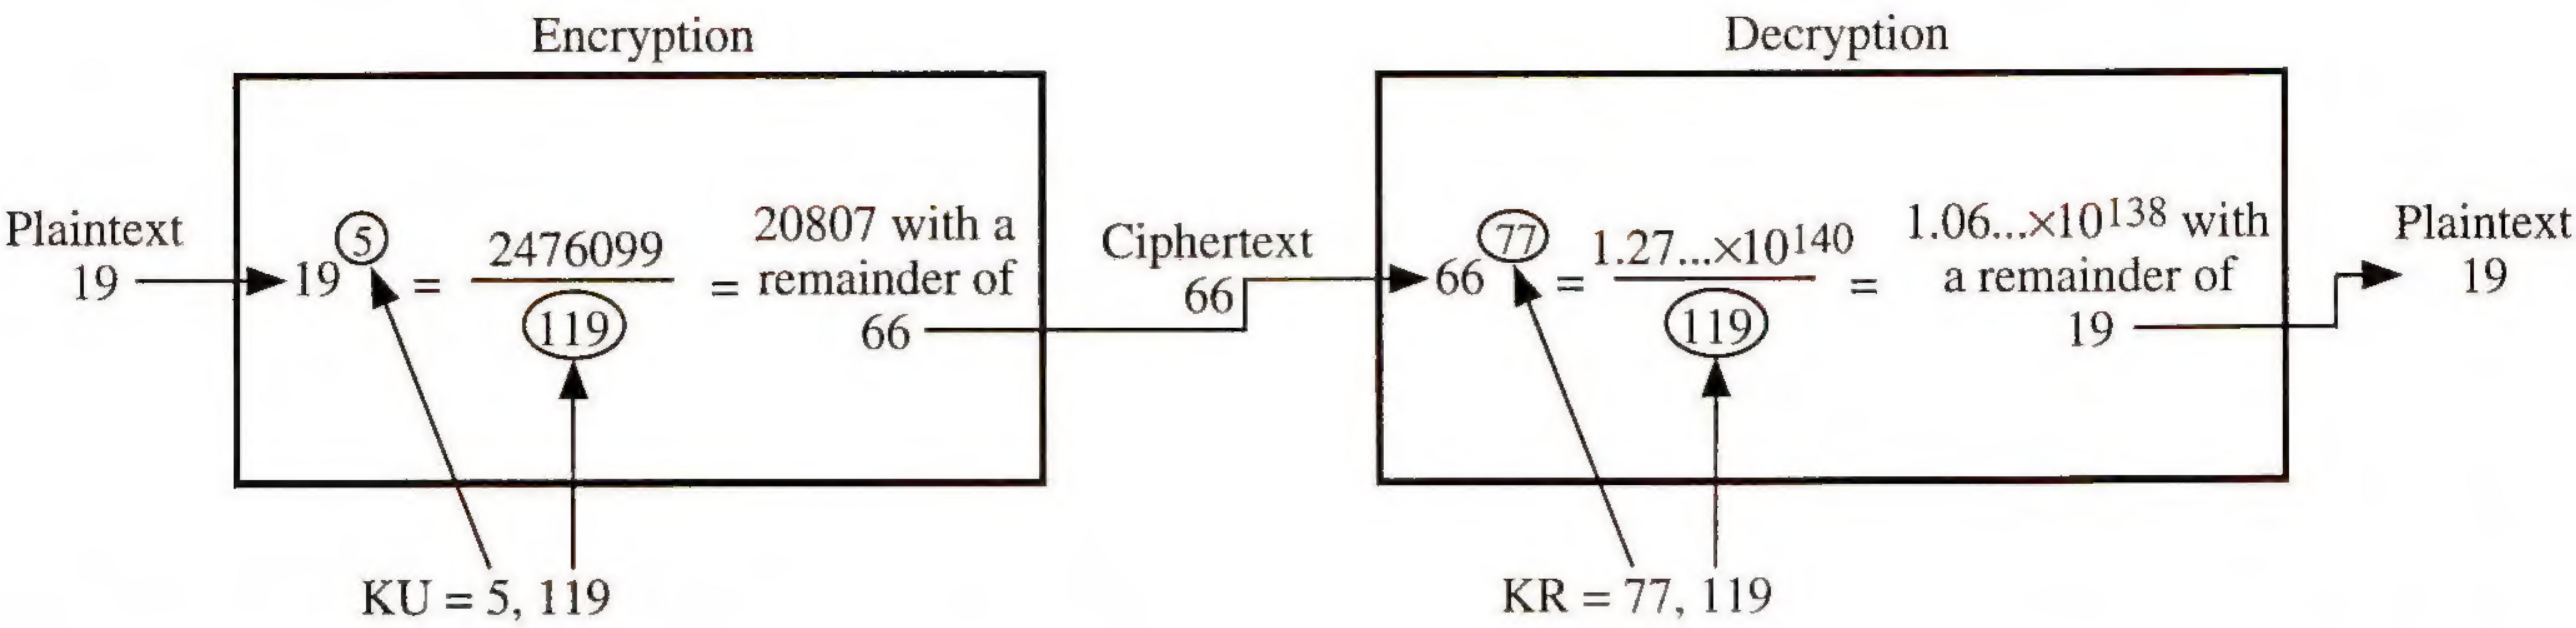
Cryptographic Algorithms *(continued)*

Figure 3: Example of RSA Algorithm

There are two possible approaches to defeating the RSA algorithm. The first is the brute force approach: try all possible private keys. Thus the larger the number of bits in e and d, the more secure the algorithm. However, because the calculations involved, both in key generation and in encryption/decryption, are complex, the larger the size of the key, the slower the system will run.

Most discussions of the cryptanalysis of RSA have focused on the task of factoring p into its two prime factors. Until recently, this was felt to be infeasible for numbers in the range of 100 decimal digits or so, which is about 300 or more bits. To demonstrate the strength of RSA, its three developers issued a challenge to decrypt a message that was encrypted using a 129-decimal-digit number as their public modulus. The authors predicted that it would take 40 quadrillion years with current technology to crack the code. Recently the code was cracked by a worldwide team cooperating over the Internet and using over 1,600 computers after only eight months of work [3]. This result does not invalidate the use of RSA; it simply means that larger key sizes must be used. Currently, a 1024-bit key size (about 300 decimal digits), is considered strong enough for virtually all applications.

Digital Signature Standard (DSS)

There is one other public-key scheme worthy of mention (Table 1). The National Institute of Standards and Technology (NIST) has published a Federal Information Processing Standard known as the *Digital Signature Standard* (DSS) [4]. The DSS was originally proposed in 1991 and revised in 1993 in response to public feedback concerning the security of the scheme.

Algorithm	Functionality	Example applications used in
RSA	Digital signature; encryption	PEM, PGP
DSA	Digital signature	DSS

RSA = Rivest, Shamir, Adleman algorithm
DSA = Digital Signature Algorithm
PEM = Privacy Enhanced Mail
PGP = Pretty Good Privacy
DSS = Digital Signature Standard

Table 1: Noteworthy Public-Key Cryptographic Algorithms

The DSS uses an algorithm that is designed to provide only the digital signature function. Unlike RSA, it cannot be used for encryption. Nevertheless, it is a public-key technique. The algorithm itself is quite complex and will not be described here.

It is worth noting that NIST asserts that the DSS is free of any patent claims. However, Public Key Partners, Inc. (PKP), holder of a number of encryption patents including that for RSA, claims that DSS infringes on some of these patents. The matter has yet to be tested in court.

Secure Hash Functions

A hash value is generated by a function H of the form:

$$h = H(M)$$

where M is a variable-length message, and H(M) is the fixed-length hash value. The hash value is appended to the message at the source at a time when the message is assumed or known to be correct. The receiver authenticates that message by recomputing the hash value. Because the hash function itself is not considered to be secret, some means is required to protect the hash value.

The purpose of a hash function is to produce a "fingerprint" of a file, message, or other block of data. There are two common applications for hash functions: message authentication and digital signature. One way to use a hash function for authentication is to append a secret value shared by transmitter and receiver to a message and take the hash function of the message plus secret value. Then the message and hash function are transmitted. At the receiver, the secret value is again appended to the message and the hash function recomputed. If the two hash values match, it is assumed that the message has not been altered.

For digital signature applications, the source computes the hash function of a message and then encrypts the hash function with its private key. The message plus encrypted hash function are then transmitted. The receiver can again calculate the hash function and can recover the transmitted hash function by decrypting with the sender's public key. Again, a match assures that the message has not been altered. Furthermore, the encrypted hash function acts as a signature since only the alleged source knows the required private key.

To be useful for message authentication and digital signature, a hash function H must have the following properties:

1. H can be applied to a block of data of any size.

2. H produces a fixed-length output.

3. $H(x)$ is relatively easy to compute for any given x, making both hardware and software implementations practical.

4. For any given code m, it is computationally infeasible to find x such that $H(x) = m$.

5. For any given block x, it is computationally infeasible to find $y \neq x$ with $H(y) = H(x)$.

6. It is computationally infeasible to find any pair (x, y) such that $H(x) = H(y)$.

The first three properties are requirements for the practical application of a hash function to message authentication.

The fourth property is the "one-way" property: It is easy to generate a code given a message but virtually impossible to generate a message given a code. This property is important if the authentication technique involves the use of a secret value.

Cryptographic Algorithms *(continued)*

The secret value itself is not sent; however, if the hash function is not one-way, an attacker can easily discover the secret value: If the attacker can observe or intercept a transmission, the attacker obtains the message M and the hash code $C = H(S_{AB}\|M)$. The attacker then inverts the hash function to obtain $S_{AB}\|M = H^{-1}(C)$. Because the attacker now has both M and $S_{AB}\|M$, it is a trivial matter to recover S_{AB}.

The fifth property guarantees that an alternative message hashing to the same value as a given message cannot be found. This prevents forgery when an encrypted hash code is used. For these cases, the opponent can read the message and therefore generate its hash code. But, because the opponent does not have the secret key, the opponent should not be able to alter the message without detection. If this property were not true, an attacker would be capable of the following sequence: First, observe or intercept a message plus its encrypted hash code; second, generate an unencrypted hash code from the message; third, generate an alternate message with the same hash code.

A hash function that satisfies the first five properties in the preceding list is referred to as a *weak* hash function. If the sixth property is also satisfied, then it is referred to as a *strong* hash function. The sixth property protects against a sophisticated class of attack known as the "Birthday Attack." Two of the most secure hash functions are listed in Table 2.

Algorithm	Hash code size (bits)	Example applications used in
MD5	128	PEM, PGP, SNMPv2
SHA	160	DSS

MD5 = Message Digest, version 5
SHA = Secure Hash Algorithm
SNMPv2 = Simple Network Management Protocol, version 2
PGP = Pretty Good Privacy
PEM = Privacy Enhanced Mail
DSS = Digital Signature Standard

Table 2: Noteworthy Secure Hash Algorithms

MD5 The MD5 message-digest algorithm [5] was developed by Ron Rivest at MIT (the "R" in the RSA [Rivest-Shamir-Adelman] public-key encryption algorithm). The algorithm takes as input a message of arbitrary length and produces as output a 128-bit message digest. The input is processed in 512-bit blocks.

Figure 4 depicts the overall processing of a message to produce a digest. The processing consists of the following steps:

• *Step 1: Append Padding Bits:* The message is padded so that its length in bits is congruent to 448 modulo 512 (length ≡ 448 mod 512). That is, the length of the padded message is 64 bits less than an integer multiple of 512 bits. Padding is always added, even if the message is already of the desired length. For example, if the message is 448 bits long, it is padded by 512 bits to a length of 960 bits. Thus, the number of padding bits is in the range of 1 to 512. The padding consists of a single 1-bit followed by the necessary number of 0-bits.

• *Step 2: Append Length:* A 64-bit representation of the length in bits of the original message (before the padding) is appended to the result of step 1. If the original length is greater than 2^{64}, then only the low-order 64 bits of the length are used. Thus, the field contains the length of the original message, modulo 2^{64}.

The outcome of the first two steps yields a message that is an integer multiple of 512 bits in length. In the figure, the expanded message is represented as the sequence of 512-bit blocks Y_0, Y_1, . . ., Y_{L-1}, so that the total length of the expanded message is $L \times 512$ bits. Equivalently, the result is a multiple of 16 32-bit words. Let M[0 . . . N–1] denote the words of the resulting message, with N an integer multiple of 16. Thus $N = L \times 16$.
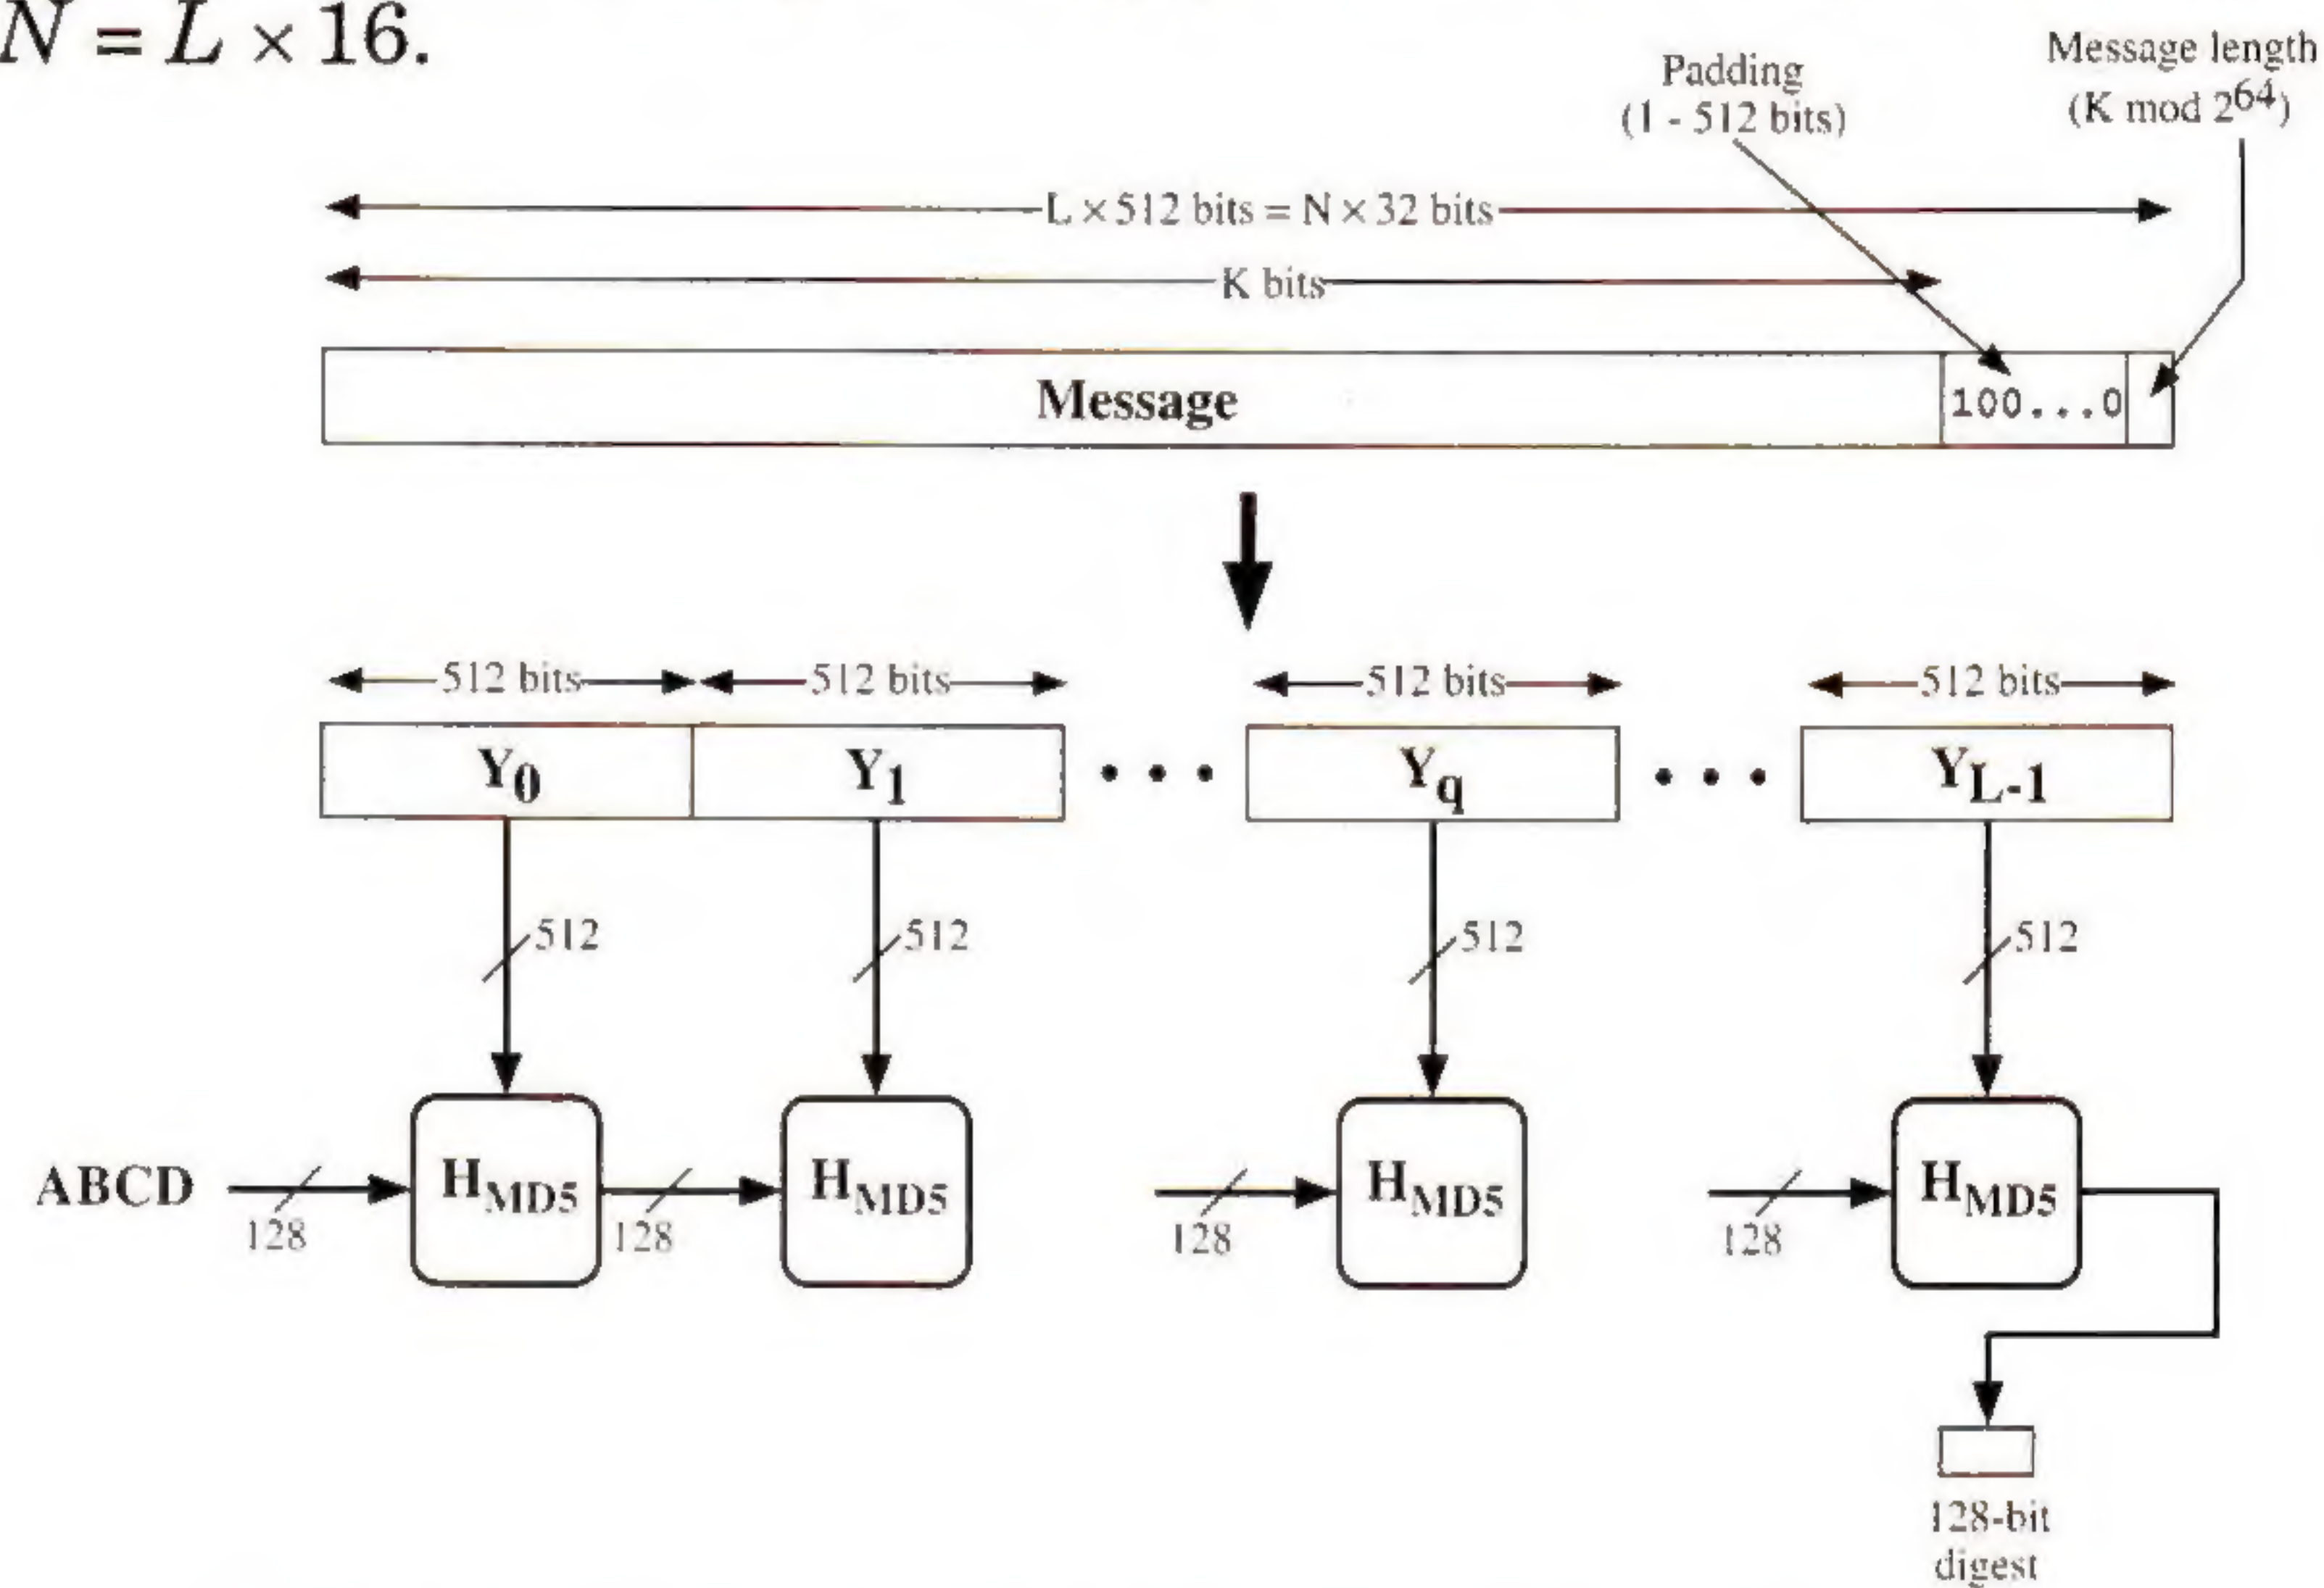

Figure 4: Message Digest Generation Using MD5

• *Step 3: Initialize MD Buffer:* A 128-bit buffer is used to hold intermediate and final results of the hash function. The buffer can be represented as four 32-bit registers (A, B, C, D). These registers are initialized to the following hexadecimal values (low-order octets first):

```
A = 01234567
B = 89ABCDEF
C = FEDCBA98
D = 76543210
```

• *Step 4: Process Message in 512-Bit (16-Word) Blocks:* The heart of the algorithm is a module that consists of four "rounds" of processing; this module is labeled H_{MD5} in Figure 4, and its logic is illustrated in Figure 5. The four rounds have a similar structure but each uses a different primitive logical function, referred to as F, G, H, and I in the specification. In the figure, the four rounds are labeled f_F, f_G, f_H, f_I, to indicate that each round has the same general functional structure, f, but depends on a different primitive function (F, G, H, I).

Note that each round takes as input the current 512-bit block being processed (Y_q) and the 128-bit buffer value ABCD and updates the contents of the buffer. Each round also makes use of one-fourth of a 64-element table T[1 . . . 64], constructed from the sine function.

The ith element of T, denoted T[i], has the value equal to the integer part of $2^{32} \times$ abs[(sin(i)], where i is in radians. Since abs[(sin(i)] is a number between 0 and 1, each element of T is an integer that can be represented in 32 bits. The table provides a "randomized" set of 32-bit patterns, which should eliminate any regularities in the input data.
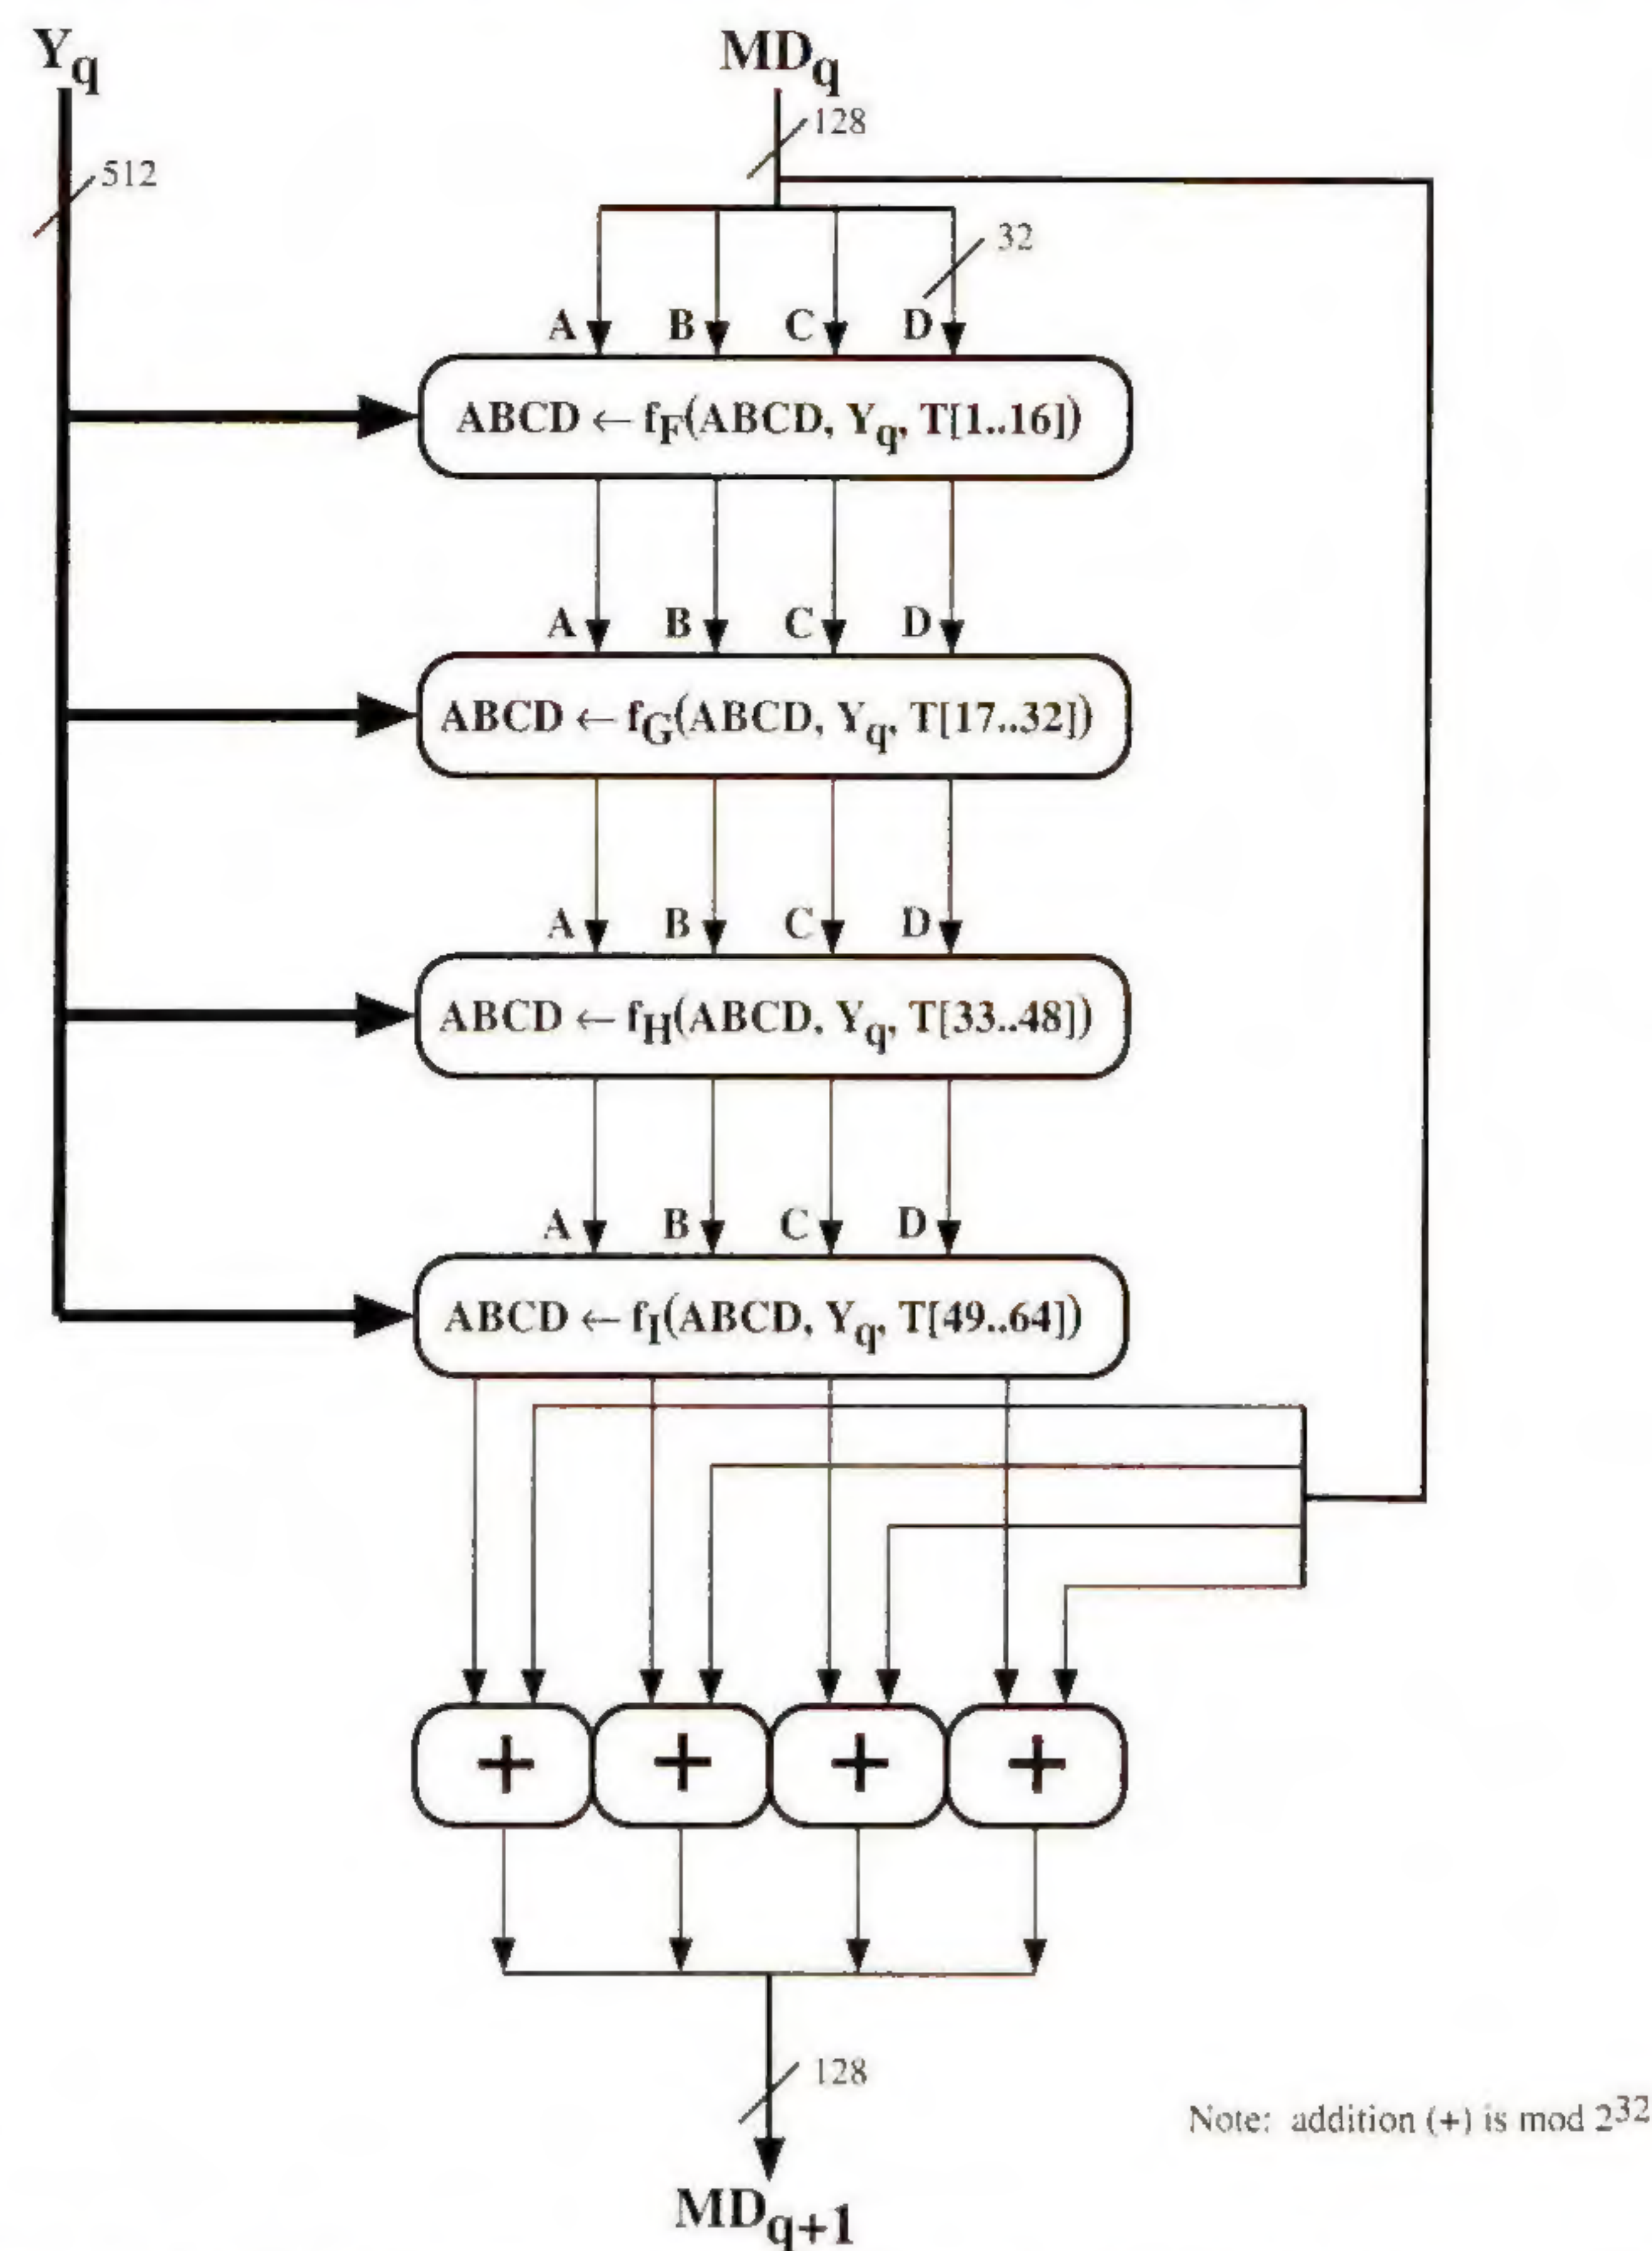
9

Cryptographic Algorithms *(continued)*

Figure 5: MD5 Processing of Single 512-bit Block (HD$_{MD5}$)

Overall, for block Y$_q$, the algorithm takes Y$_q$ and an intermediate digest value MD$_q$ as inputs. MD$_q$ is placed into buffer ABCD. The output of the fourth round is added to MD$_q$ to produce MD$_{q+1}$. The addition is done independently for each of the four words in the buffer with each of the corresponding words in MD$_q$, using addition modulo 2^{32}.

• *Step 5: Output*: After all *L* 512-bit blocks have been processed, the output from the Lth stage is the 128-bit message digest.

Each of the functions labeled, f$_F$, f$_G$, f$_H$, f$_I$, involves bitwise logical AND, OR, and XOR operations as well as addition modulo 2^{32}.

The MD5 algorithm has the property that every bit of the hash code is a function of every bit in the input. The complex repetition of the basic functions (F, G, H, I) produces results that are well mixed; that is, it is unlikely that two messages chosen at random, even if they exhibit similar regularities, will have the same hash code. Rivest conjectures in the RFC that MD5 is as strong as possible for a 128-bit hash code; namely, the difficulty of coming up with two messages having the same message digest is on the order of 2^{64} operations, whereas the difficulty of finding a message with a given digest is on the order of 2^{128} operations. As of this writing, no analysis has been done to disprove these conjectures.

Secure Hash Algorithm (SHA)

The *Secure Hash Algorithm* (SHA) was developed by the National Institute of Standards and Technology (NIST) and published as a federal information processing standard in 1993 [6]. SHA is based on the MD4 algorithm, the precursor of MD5, and its design closely models MD4. SHA is used as part of the Digital Signature Standard, but can be used in any security application that requires a hash code.

WILLIAM STALLINGS is an independent consultant whose clients have included major corporations and government agencies in the United States and Europe. He is the author over over a dozen books on data communications and computers, including *Data and Computer Communications, Fourth Edition,* from Prentice-Hall. He is currently at work on *Protect Your Privacy: The User's Guide to PGP,* to be published by Prentice Hall in November. He holds a PhD from M.I.T. in Computer Science and a B.S. from Notre Dame in electrical engineering. E-mail: `stallings@acm.org`

	MD5	SHA
Digest length	128 bits	160 bits
Basic unit of processing	512 bits	512 bits
Number of steps	64 (four rounds of 16)	80
Maximum message size	∞	2^{64} bits
Primitive logical functions	4	3
Additive constants used	64	4

Table 3: Comparison of MD5 and SHA

Table 3 summarizes key differences between the two algorithms. We can make the following observations:

• *Security:* The most obvious, and most important, difference, is that the SHA digest is 32 bits longer than the MD5 digest. If neither algorithm contains any structural flaws that are vulnerable to cryptanalytic attack, which is so far a reasonable assumption, then SHA is the stronger algorithm. Using a brute-force technique, the difficulty of producing any message having a given message digest is on the order of 2^{128} operations for MD5 and 2^{160} for SHA. Again, using a brute-force technique, the difficulty of producing two messages having the same message digest is on the order of 2^{64} operations for MD5 and 2^{80} for SHA.

• *Speed:* Because both algorithms rely heavily on addition modulo 2^{32}, both do well on a 32-bit architecture. SHA involves more steps (80 versus 64) and must process a 160-bit buffer compared to MD5's 128-bit buffer. Thus, SHA should execute about 25% slower than MD5 on the same hardware.

• *Simplicity and compactness:* Both algorithm are simple to describe and simple to implement, without requiring large programs or substitution tables. However, SHA uses a single step structure, compared to the four structures used in MD5. Furthermore, the manipulation of the buffer words is the same for all SHA steps, whereas in MD5, the arrangement of words is specified individually for each step. Thus, in this category, SHA gets the nod.

A final note: On April 22, 1994, NIST issued an advisory that there would be a technical revision to SHA. The revision is to correct a minor flaw discovered in the algorithm. NIST claims that the original SHA remains highly secure but that the revision will increase that security.

References

[1] Diffie, W., and Hellman, M., "New Directions in Cryptography," *IEEE Transactions on Information Theory,* November 1976.

[2] Rivest, R.; Shamir, A.; and Adleman, L., "A Method for Obtaining Digital Signatures and Public Key Cryptosystems," *Communications of the ACM,* February 1978.

[3] Leutwyler, K., "Superhack," *Scientific American,* July 1994.

[4] NIST, "Digital Signature Standard," FIPS PUB 186, 1994.

[5] Rivest, R., "The MD5 Message-Digest Algorithm," RFC 1321, April 1992.

[6] NIST, "Secure Hash Standard," FIPS PUB 180, 1993.

[7] Stallings, W., "Cryptographic Algorithms, Part I: Conventional Cryptography," *ConneXions,* Volume 8, No. 9, September 1994.

Interworking with B-ISDN:
Data Services and Signalling
by Reto Beeler, Ascom Tech AG

Abstract

In *Broadband ISDN* (B-ISDN), a connectionless mode is provided to support data services. In contrast to this widely accepted approach, this article proposes to support data services in a connection-oriented mode. It does not try to propose revolutionary solutions. The intention of this article is rather to revisit the current proposals having in mind possible future directions, and to put the current solutions into this new context.

An investigation of connection-oriented support of data services shows that it can be implemented with minimal additional software, as soon as signalling is introduced in B-ISDN.

Two current solutions are briefly presented for connectionless support of data services in B-ISDN. Because both of them are based on the use of semipermanent connections, they reserve the maximum required bandwidth in the network, even during low usage periods. This is reasonable for early cross-connect networks, but wasteful in networks supporting signalling. It is shown that practical user-friendly implementations of signalling should feature an automatic signalling activation to hide the technical details to the user. By introducing the automatic signalling activation, there are nearly all the necessary elements for connection-oriented support of data services. On top of the signalling activation layer any network layer may be used, and any network can be interworked with ATM via its network layer and the corresponding signalling activation layer. As an example of this concept, it is shown how a signalling activation layer for IP based data services can be implemented. A final comparison with the connectionless mode of B-ISDN reveals that this is presumably to be replaced by solutions with automatic signalling activation, as soon as signalling is introduced in B-ISDN.

Introduction

It is generally accepted by researchers and marketing people that data services will contribute substantially to the load of early B-ISDN. Therefore, the support of connectionless services has been discussed extensively in the relevant standard bodies, and in the research community. The reason for this huge interest in connectionless services in B-ISDN is based on the claims that current data networks are of a connectionless nature, and that data traffic will generate a substantial fraction of the load in early B-ISDN. In current data networks, the data are sent as soon as the media access control grants permission; there is no connection set-up nor connection release. It has been assumed for a long time that B-ISDN has to accept such a spontaneous traffic, and that there is no time to invoke a channel set-up, especially at high bit rates. This led to particular solutions like the *Switched Multimegabit Data Service* (SMDS) and the *Connectionless Broadband Data Service* (CBDS). Today, there are some practical solutions on the marketplace which feature a proprietary signalling protocol being activated by IP traffic [1]. In addition to that, current discussions on LAN-emulation over ATM show that an SMDS type service is not the only answer to the data traffic problem in B-ISDN. The approach outlined in this article highlights the problem from an application based viewpoint. It tries to generalise the approach of automatic signalling activation and claims that some assumptions on data traffic in B-ISDN should be reconsidered.

Current support for data services in B-ISDN

Data services in a public B-ISDN are currently supported by SMDS [2] or related connectionless data services like CBDS [3]. The ITU-TS is currently defining an international recommendation for this kind of approach [4]. All the mentioned approaches provide a connectionless service by a network layer stacked on top of the *ATM Adaptation Layer* (AAL) type 3/4. The network protocol used in the access network is called *Connectionless Network Access Protocol* (CLNAP). Furthermore, the connectionless service can be supported by servers in the network, as shown in Figure 1. The basic service provided by this public connectionless network is packet routing. Because of the lack of signalling, the packets have to contain all information necessary for this routing. The Source and Destination addresses contained in such a packet use the ITU-T E.164 format that is also the address format for N-ISDN. In addition to packet routing, the network provides a set of access functions. SMDS/CBDS offers to the user a rather limited set of properties that nevertheless comply with the most important requirements of business applications:

- Security measures (using address validation)

- Closed user groups/virtual private networks (using address screening functions)

- Data throughput monitoring (using access class enforcement)

- Charging on a per volume basis

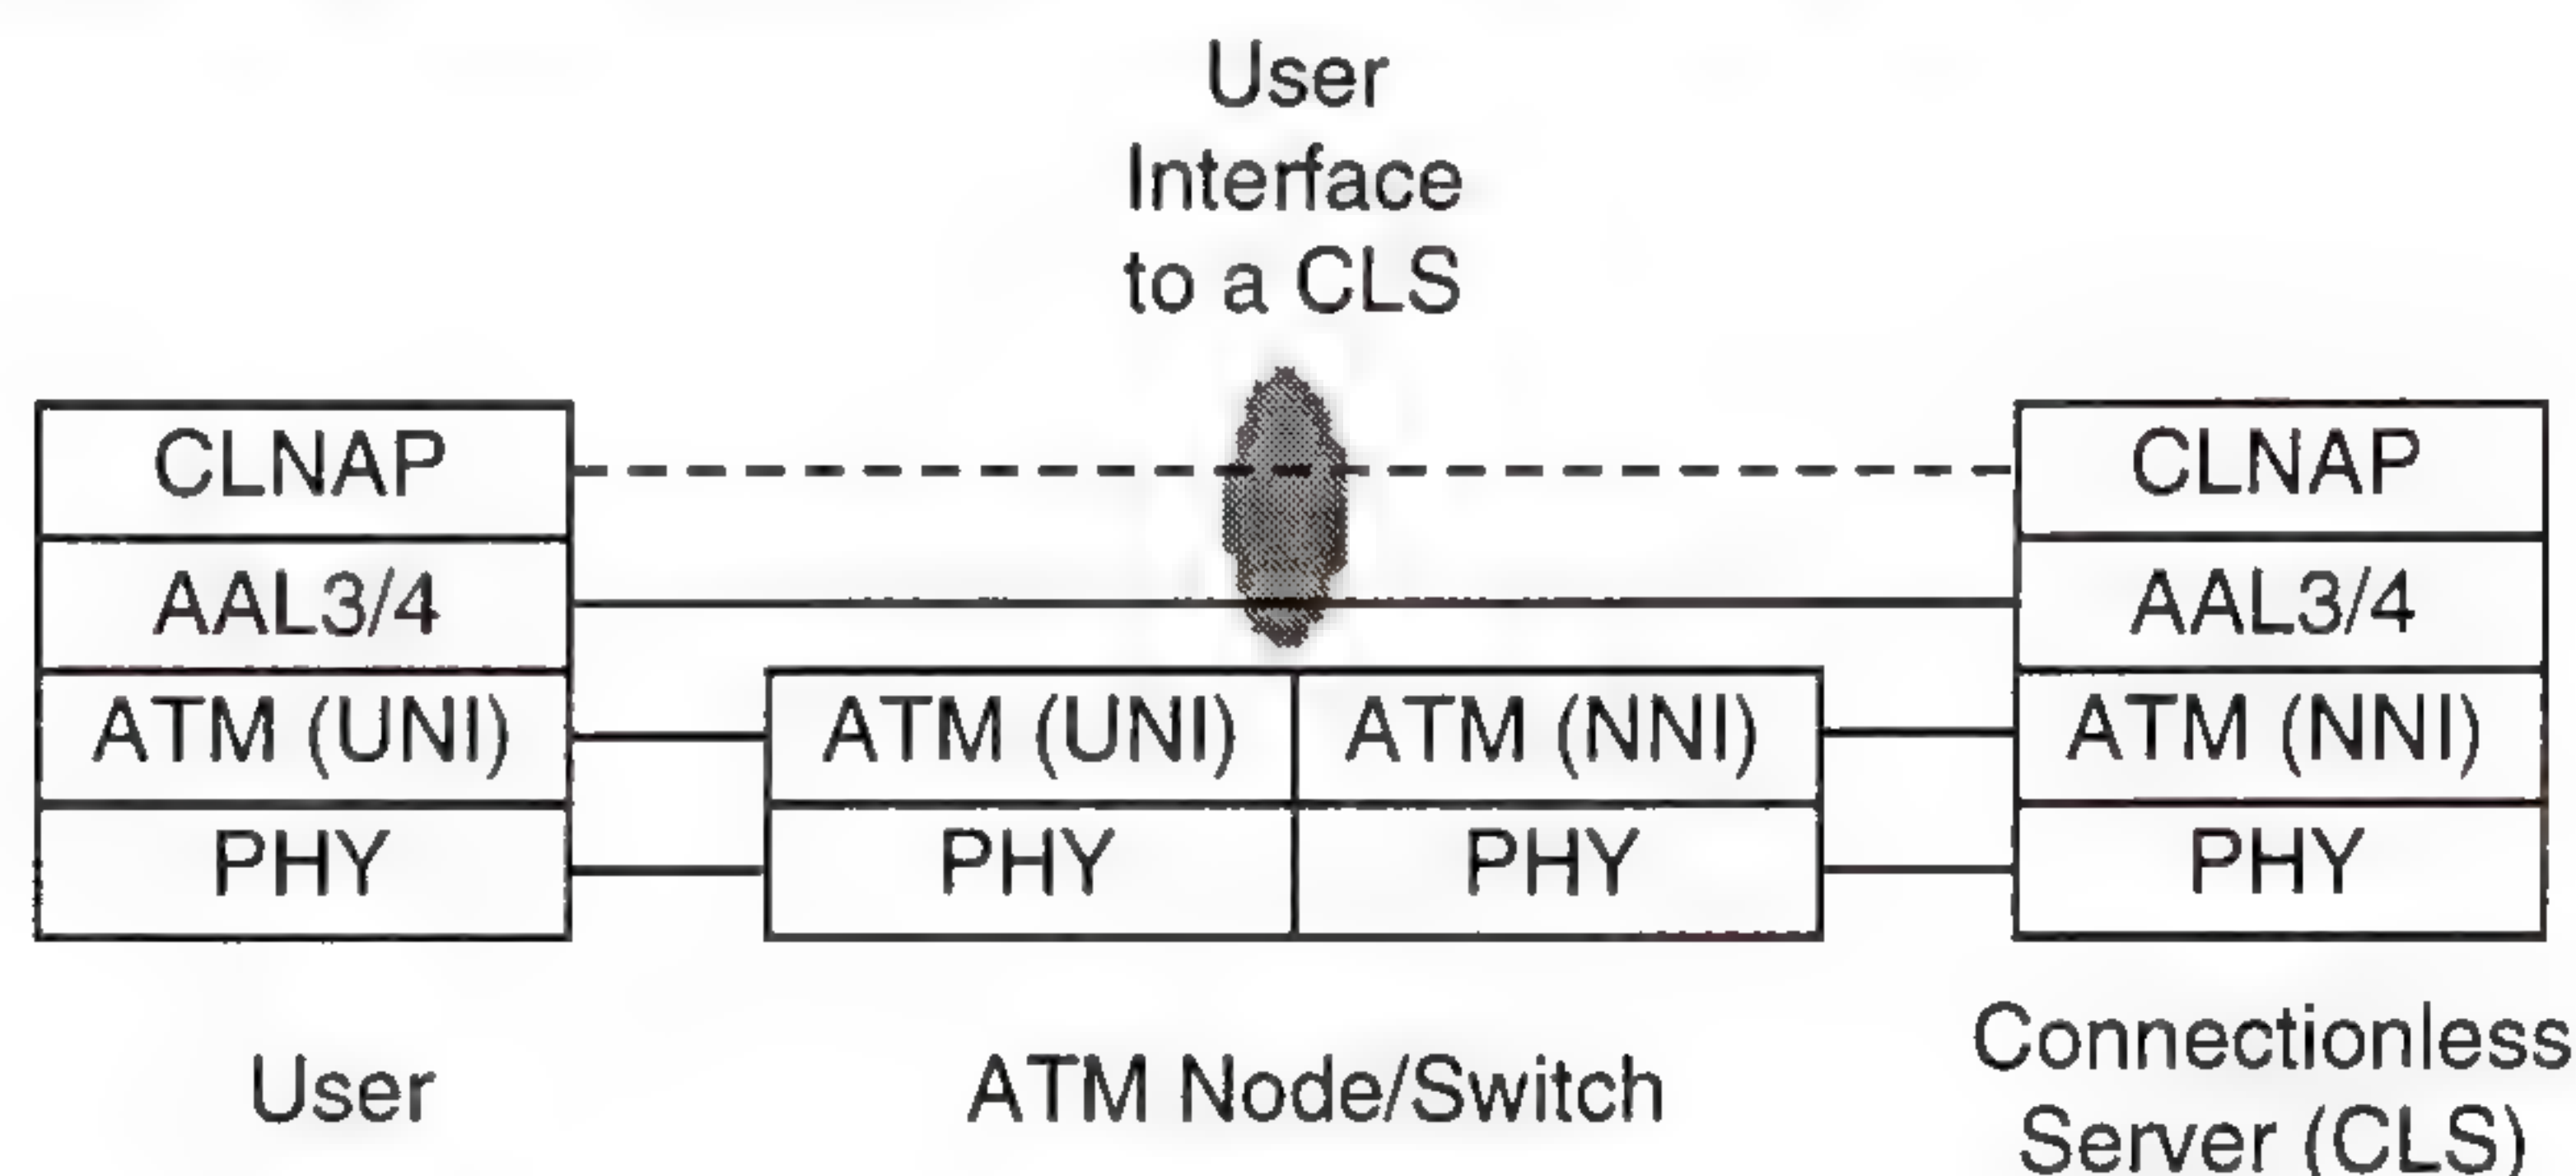- Multicast possibility (using group addressing)

Figure 1. Connectionless Service Support in B-ISDN

Virtual connections are not dynamically set-up through the B-ISDN. Instead of a dynamic path allocation there are semipermanent connections with assigned *Quality of Service* (QoS) parameters. There are two methods for such a provisioning of connectionless services in B-ISDN: the direct method and the indirect method. With the direct method, each access point offering a SMDS/CBDS type service is connected to a *Connectionless Server* (CLS). Different Connectionless Servers in the B-ISDN may be interconnected to a mesh overlay network. This approach is trying to optimise the usage of semipermanent channels, but it requires servers in the network, which process information up to the network layer. With the indirect method, each access point offering a SMDS/CBDS service is connected with all other existing access points by semi-permanent connections, i.e., the access points have to offer all the functionality required for connectionless services themselves. This approach requires a semi-permanent channel for each potential interconnection of networks attached to the B-ISDN, and therefore is quite wasteful in network resources.

13

Interworking with B-ISDN *(continued)*

Both the direct and the indirect support for connectionless services in B-ISDN rely on the use of *ATM Adaptation Layer* (AAL) 3/4. AAL 3/4 is required especially for networks with many servers: This AAL has got a checksum on the contents of every cell, and the vital service information is packed in the first cell belonging to a CLNAP packet. This allows a pipelined implementation where each cell can be processed and forwarded independently of the other cells belonging to the same packet. This helps to meet the delay requirements of the SMDS/CBDS service. Connectionless service support is currently the only application using AAL 3/4.
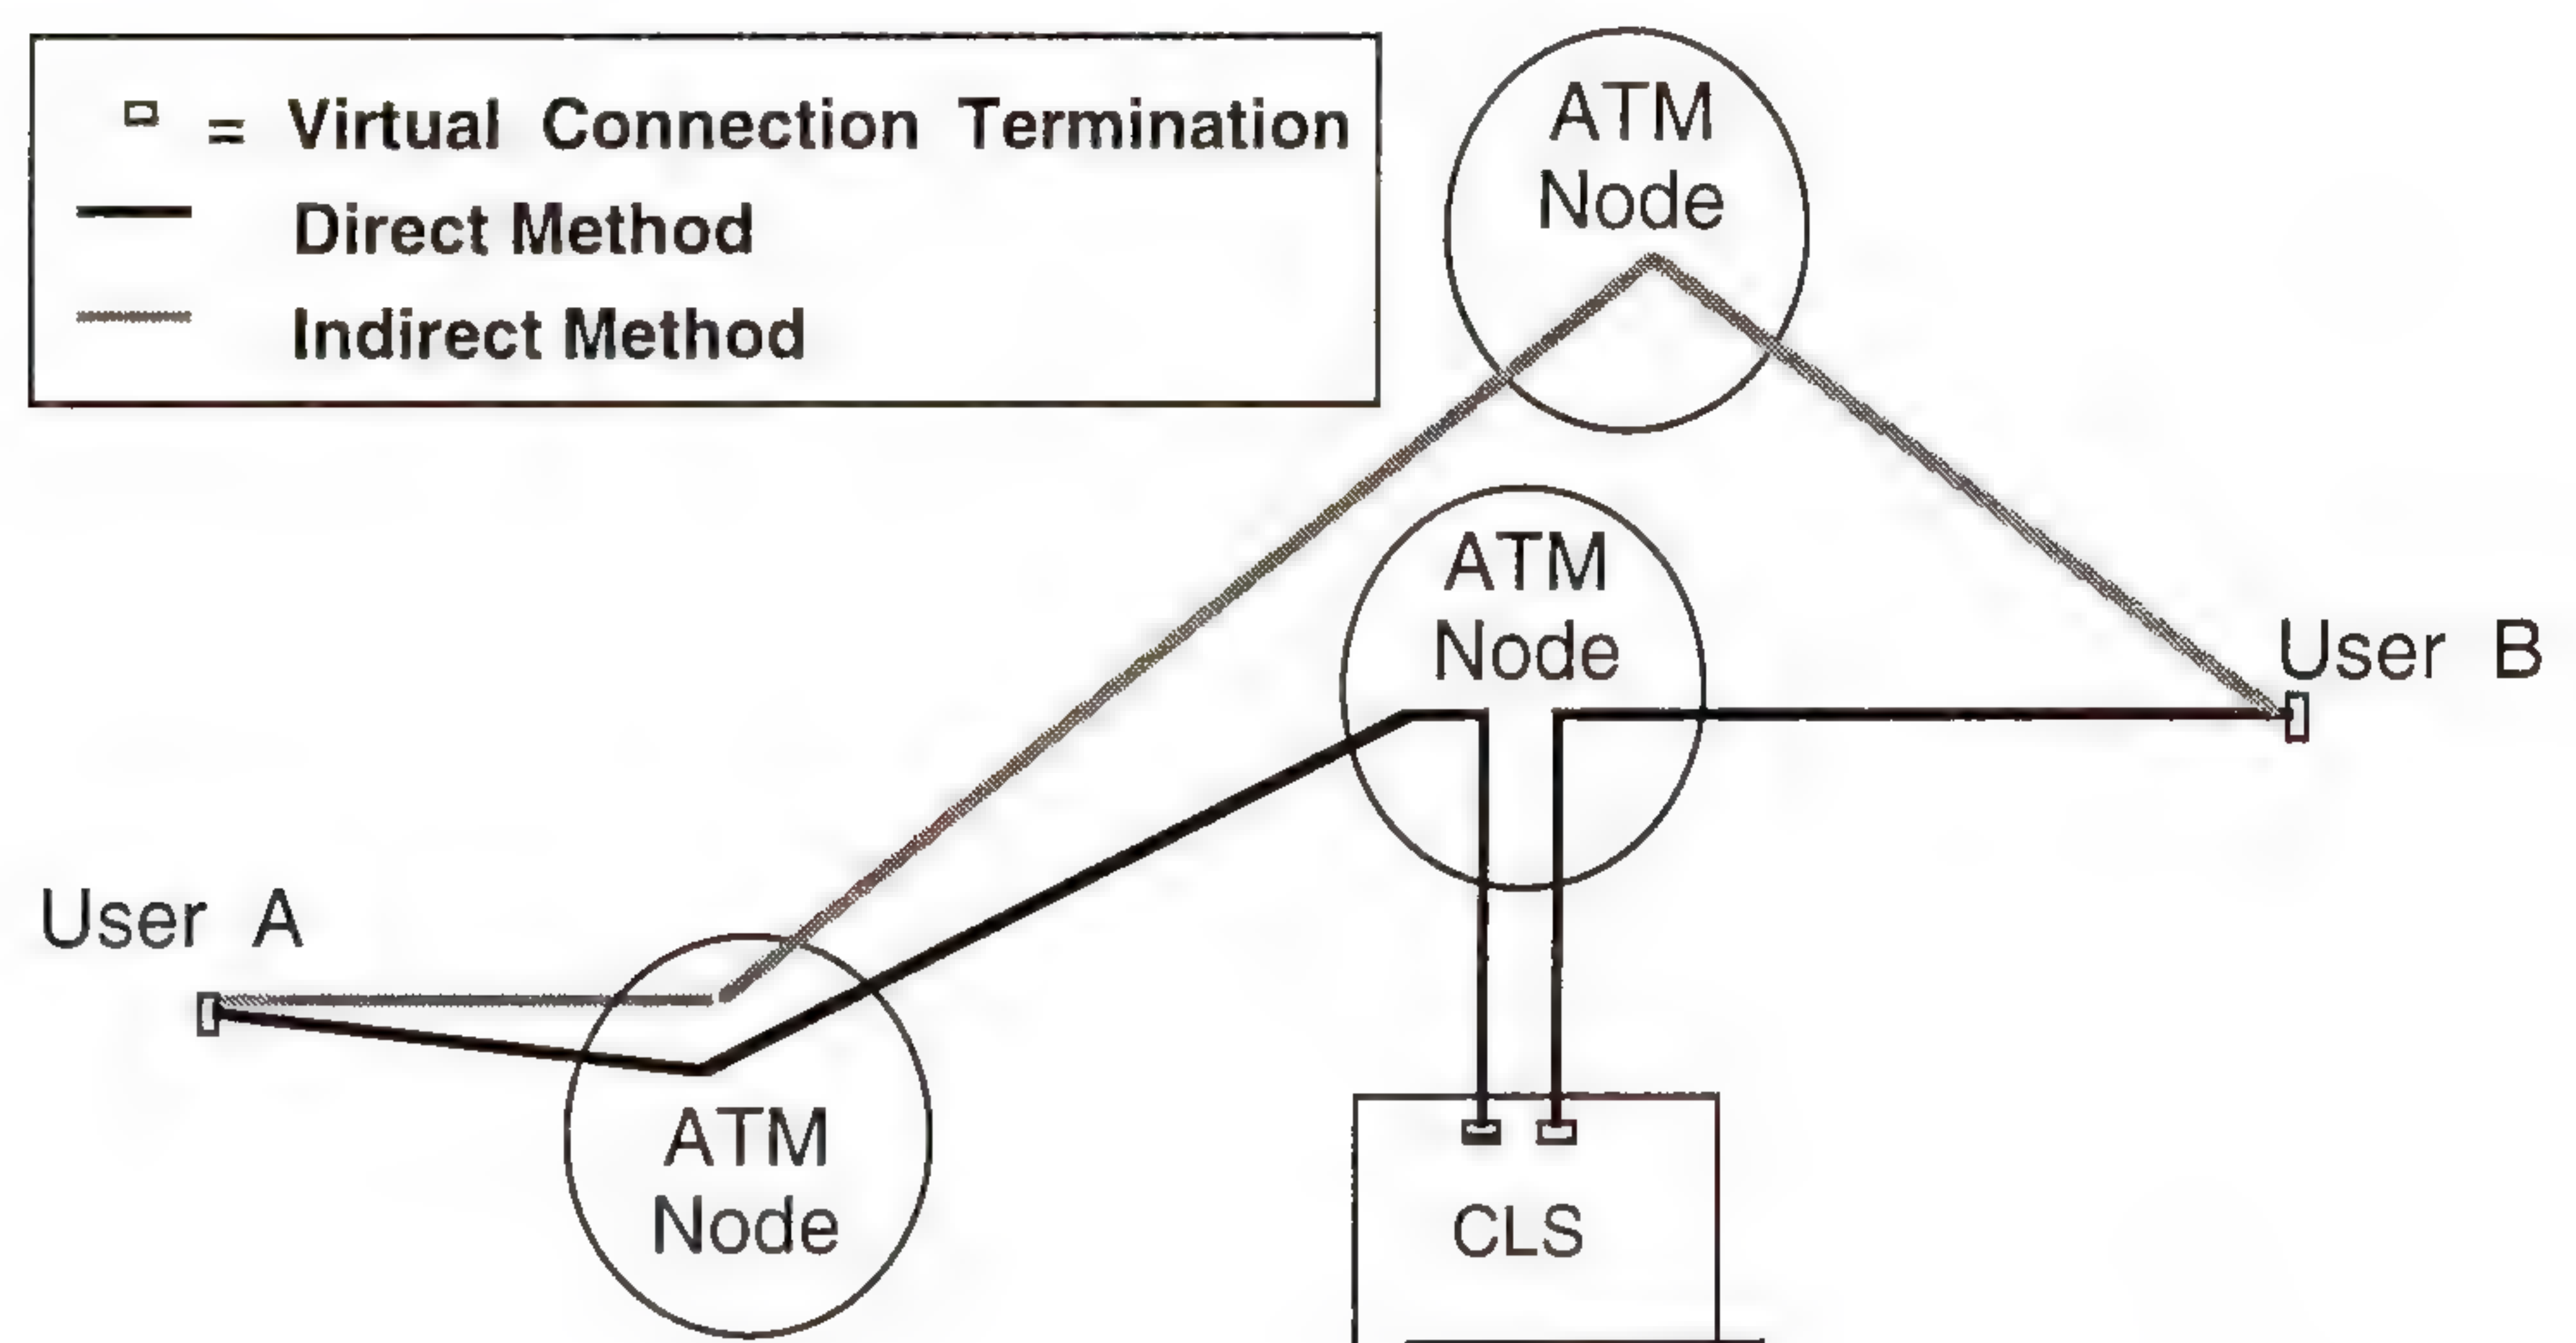

Figure 2. Direct and indirect connection methods

The argument for having a semipermanent overlay network to provide connectionless services in B-ISDN is that current data networks are of connectionless nature. This means that data are sent as soon as the media access protocol in the local area allows it. The media access control is only of local nature, and therefore the B-ISDN has to be ready to accept connectionless data at any time. This leads to the conclusion that there is no time to invoke a channel set-up by a traditional signalling protocol, especially at high bit rates. With this assumption, new signalling protocols would be required that enable a fast in-call modification of bandwidth parameters. This is a feature considered for future signalling protocols.

Because this assumption is only based on lower layer observations, it is imprecise. It is true, that the medium access control (MAC) sublayer of current LANs is connectionless. But it should be noted that by far not all data applications are connectionless. Some of them comprise a kind of connection set-up and release. The most prominent example for such applications are those based on TCP/IP, where a *three-way handshake* is used.

Furthermore, current connectionless applications do not suddenly dump very big amounts of data at high speed, because this could easily lead to data loss in Wide Area Networks. There is no guaranteed service for connectionless applications, i.e., they have to accept to use slow or lossy links. For a packet being stored during a signalling period, the memory could just be regarded as such a slow and lossy link.

From an application based point of view, the following assumptions can be taken:

- For connection-oriented applications, the set-up and release messages (e.g., three-way handshake) can be intercepted and be forwarded after the corresponding signalling has taken place.

- Connectionless data are generally designed to support loss, and they will not dump very large amounts of data at very high speed. With these assumptions, it seems to be possible to handle connections on demand.

The following sections will revisit the support of data services in B-ISDN under the assumption that some of the future B-ISDN services will use signalling. From a practical point of view, it is probable that the complicated technical details of signalling will be hidden from the user. This requires a process which shall be called *Automatic Signalling Activation* (ASA). The concept of ASA could be used for data applications as well. This leads to new conclusions for interworking data networks with B-ISDN.

Connection handling

The ATM based B-ISDN is a connection-oriented technology. Call set-up prior to user data transmission includes the request for setting up virtual channels, network resources (routes, bandwidth) and for the provision of Quality of Service (QoS) requirements. In contrast to earlier technologies, B-ISDN is more systematic as it makes a strict distinction between user plane, control plane and management plane (see Figure 3). Signalling is handled in the control plane. The ITU-T foresees to standardise three releases of signalling with the aim of backward compatibility. For each release, a corresponding capability set is defined in [5].

Release 1 is based on existing N-ISDN signalling protocols, which have been slightly adapted to handle ATM and AAL connections. Only point-to-point connections are supported, and bandwidth is allocated on a peak rate basis. Multicast connections are not yet supported. This release is now being finalised by the ITU-T in their recommendation Q.2931 [6]. A similar signalling protocol has been specified by the ATM Forum [7]. This article mainly discusses the applicability of release 1 signalling support of data services. Release 2 will introduce a new concept by separating call and bearer control. It will provide far more functionality than release 1, including a more flexible bandwidth management and complex call topologies. In-call parameter re-negotiation will add more flexibility. This might be interesting for data services, but is not extensively discussed here. Release 3 will add the functionality lacking in release 2 to fully support multimedia communications and management.

Signalling activation layer

The current discussion on connection handling in B-ISDN does not take into account the practical usage of signalling. Of course it does not, because the discussion of the control plane is restricted to the OSI network layer and below. However, for applications, the higher layers are interesting. A look at the possible provisioning of the higher layers may even introduce new aspects into interworking discussions.

B-ISDN signalling requires some generic information on the required Quality of Service (QoS). Because this information comprises technical parameters that are meaningless for a non-technical person, it can be assumed that in real-life implementations of signalling at the user side, there must be at least one additional layer in the control plane.

Interworking with B-ISDN *(continued)*

This layer, situated on top of the layer three user-to-network signalling, provides a simpler user interface. Figure 3 shows the B-ISDN protocol stack with this additional layer, named Automatic Signalling Activation (ASA).

It is, for example, very likely that a phone user on a B-ISDN network just wants to lift the handset and to dial a number. He certainly does not want to be bothered with fine tuning the required QoS parameters by setting burstiness, peak rate, mean rate etc. Therefore, these parameters are to be generated by an application dependent software process. The same reasoning applies to other services in a B-ISDN: the man-machine interface towards signalling has to be simple and straightforward. A user should only be asked to provide some basic parameters. The remaining parameters are then deduced and generated by a service specific ASA layer.
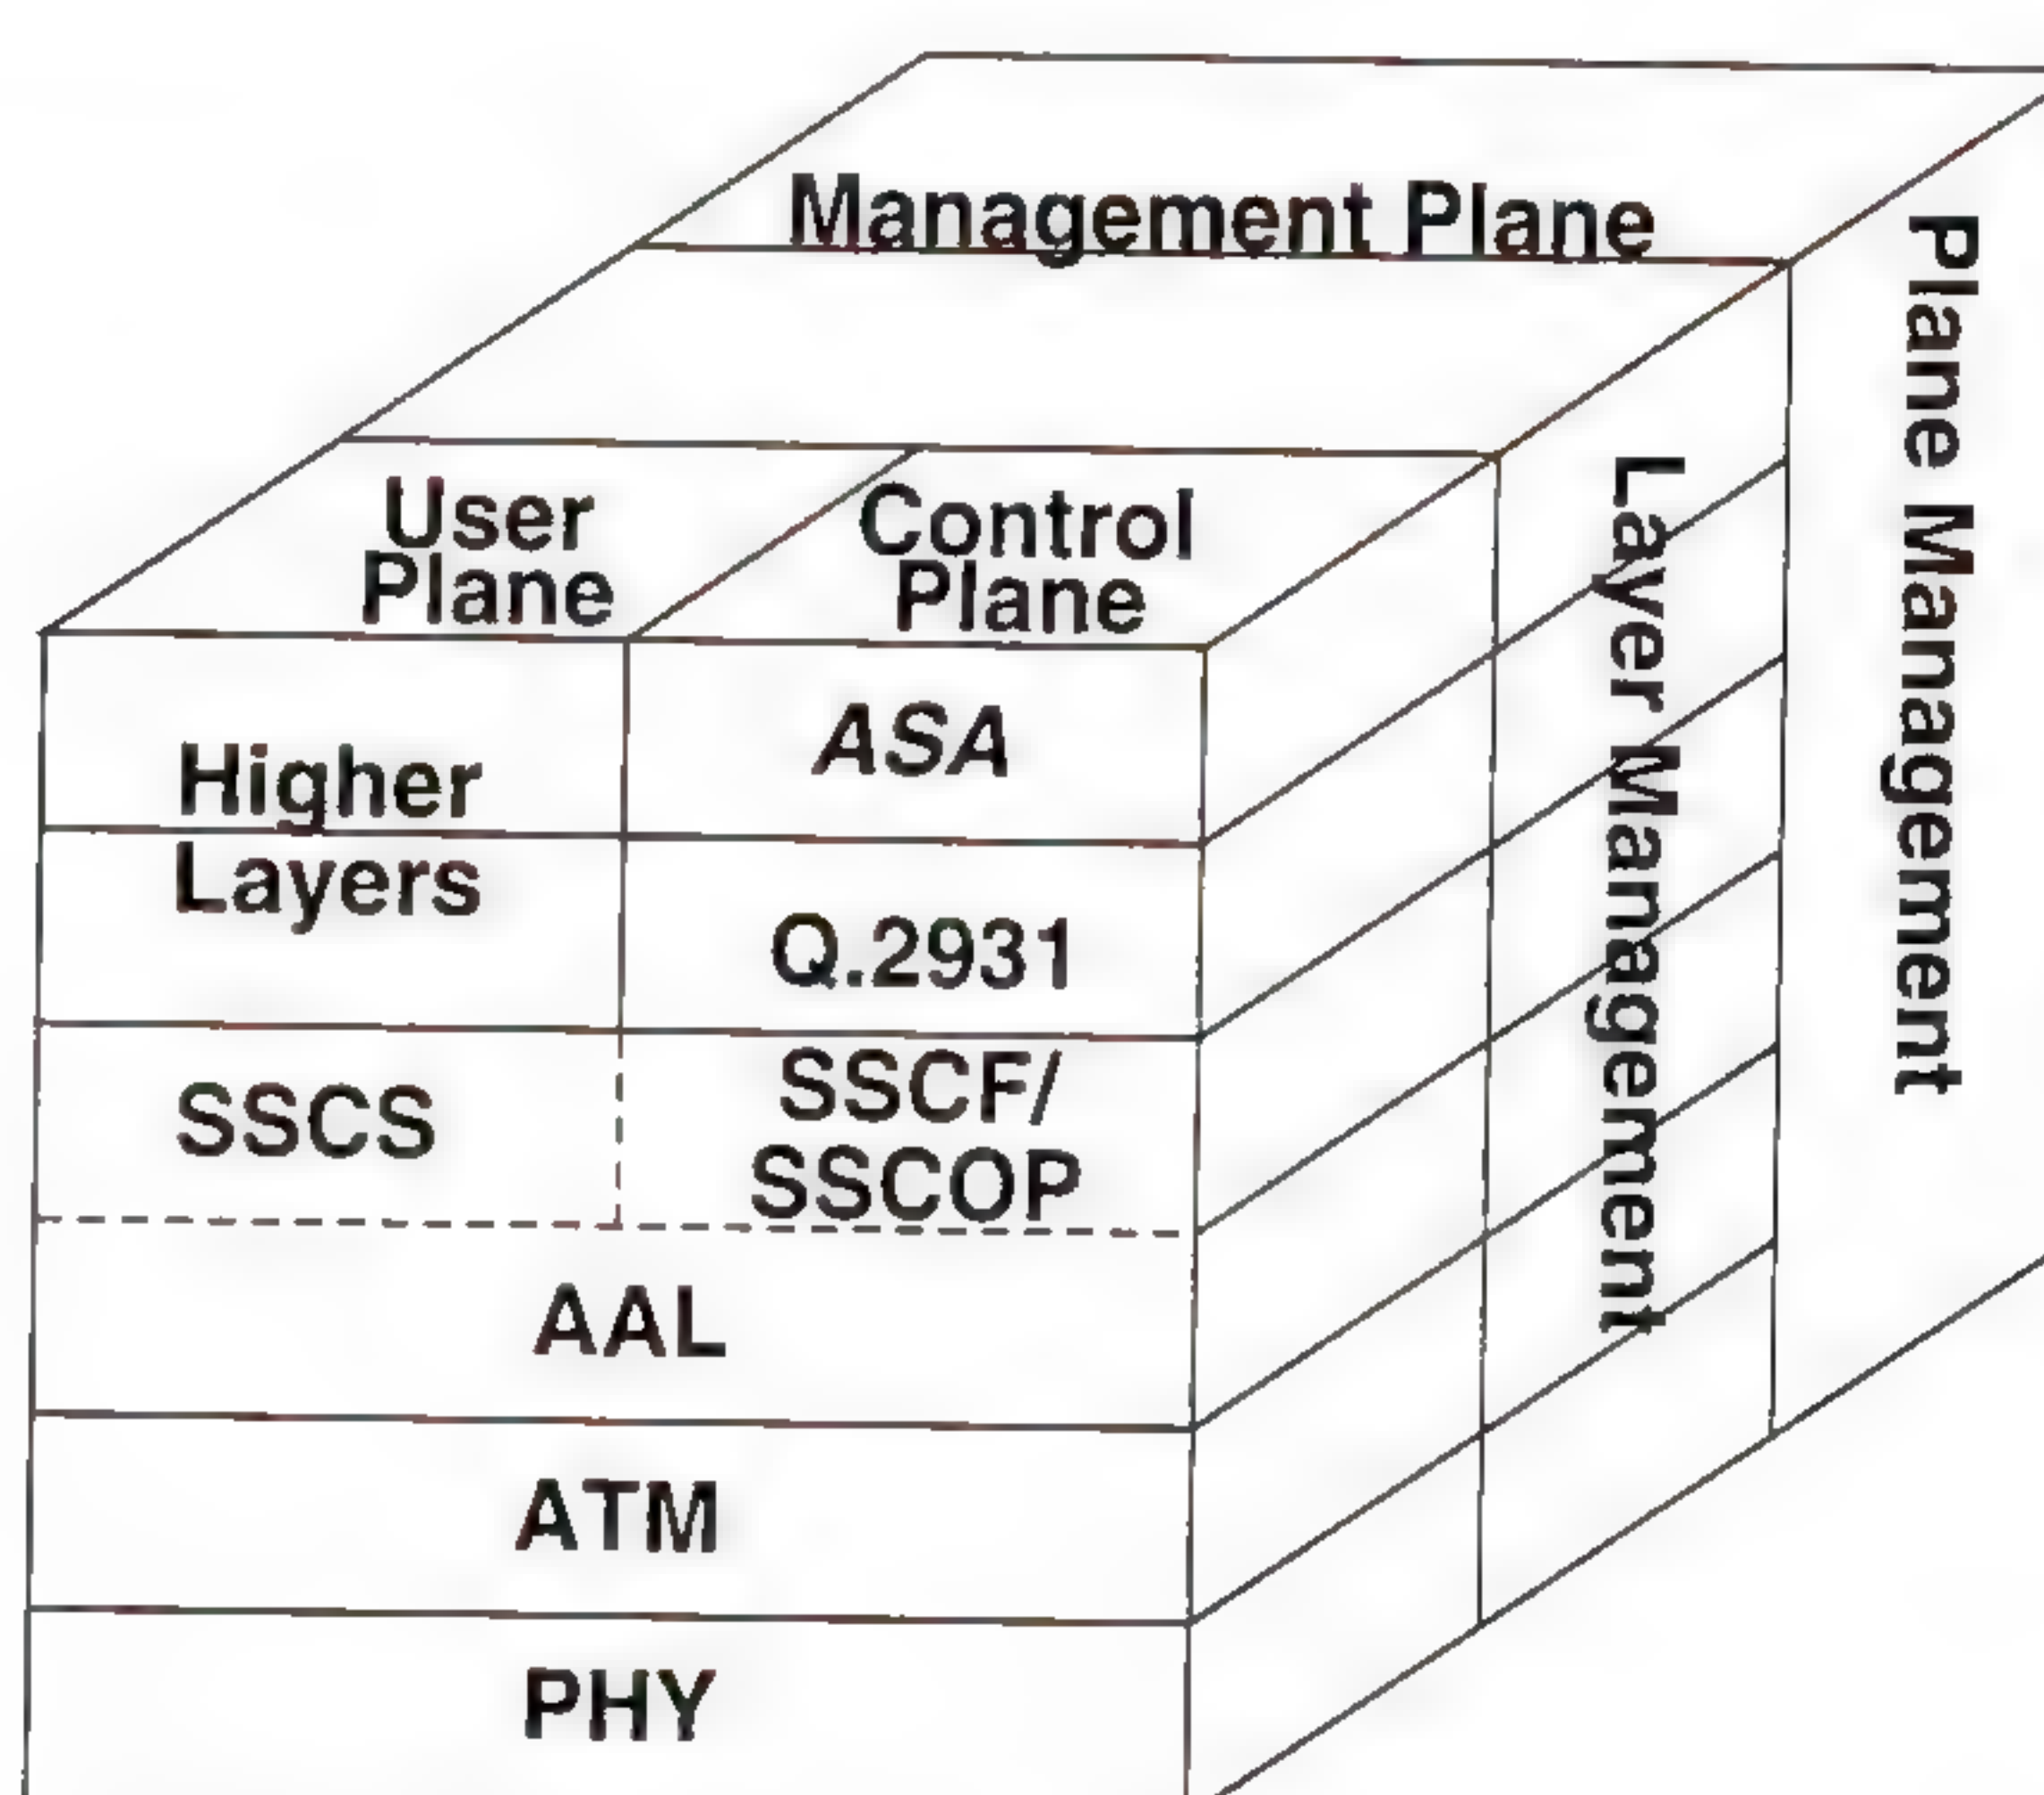

Figure 3. B-ISDN protocol stack enhanced by ASA layer

The User-Network Call Control Capabilities include the following:

- On-demand virtual channel connection set-up, maintenance and release
- Connection identifier selection/negotiation
- QoS class request/indication
- User plane AAL indication
- Calling party number services
- End-to-End transit delay
- Calling party subaddress and called party subaddress transport
- Lower-layer information
- Error handling

A user is possibly only interested in giving an identification to his or her requested service (e.g., a network address and possibly a number for additional services like calling party number services). The ASA detects the type of the connected terminal (e.g., videophone) by some kind of identification and sets up the remaining parameters automatically, possibly using a configuration database. This configuration database should be accessible via network management (based on e.g., SNMP or CMIP).

For universality it shall be assumed that a user dials a number, which is not restricted to the numbering scheme used in a public B-ISDN (E.164); in the general case the number can be viewed as an abstraction of endpoint-identification. (It could be a shortcut-number, e.g.,). In this general case, the ASA has to perform a lookup operation in the database for the network address. In principle, even a number translation between IP addresses and E.164 addresses would be possible.

Signalling for data services

It has been shown, that for practical reasons there should be an automatic signalling activation (ASA) for B-ISDN signalling. Such an ASA could even be used to support data services in B-ISDN. The problems for the ASA to support data services are the same as those for any other service: Given a network address, the signalling has to be activated. The QoS parameters have to be deduced from a configuration database for the application, which in this case is transport of data services over ATM.

As an example of this concept, a possible implementation for the support of UDP/IP and TCP/IP services is now outlined. These protocols are most likely to be responsible for a big amount of data traffic in early B-ISDN.

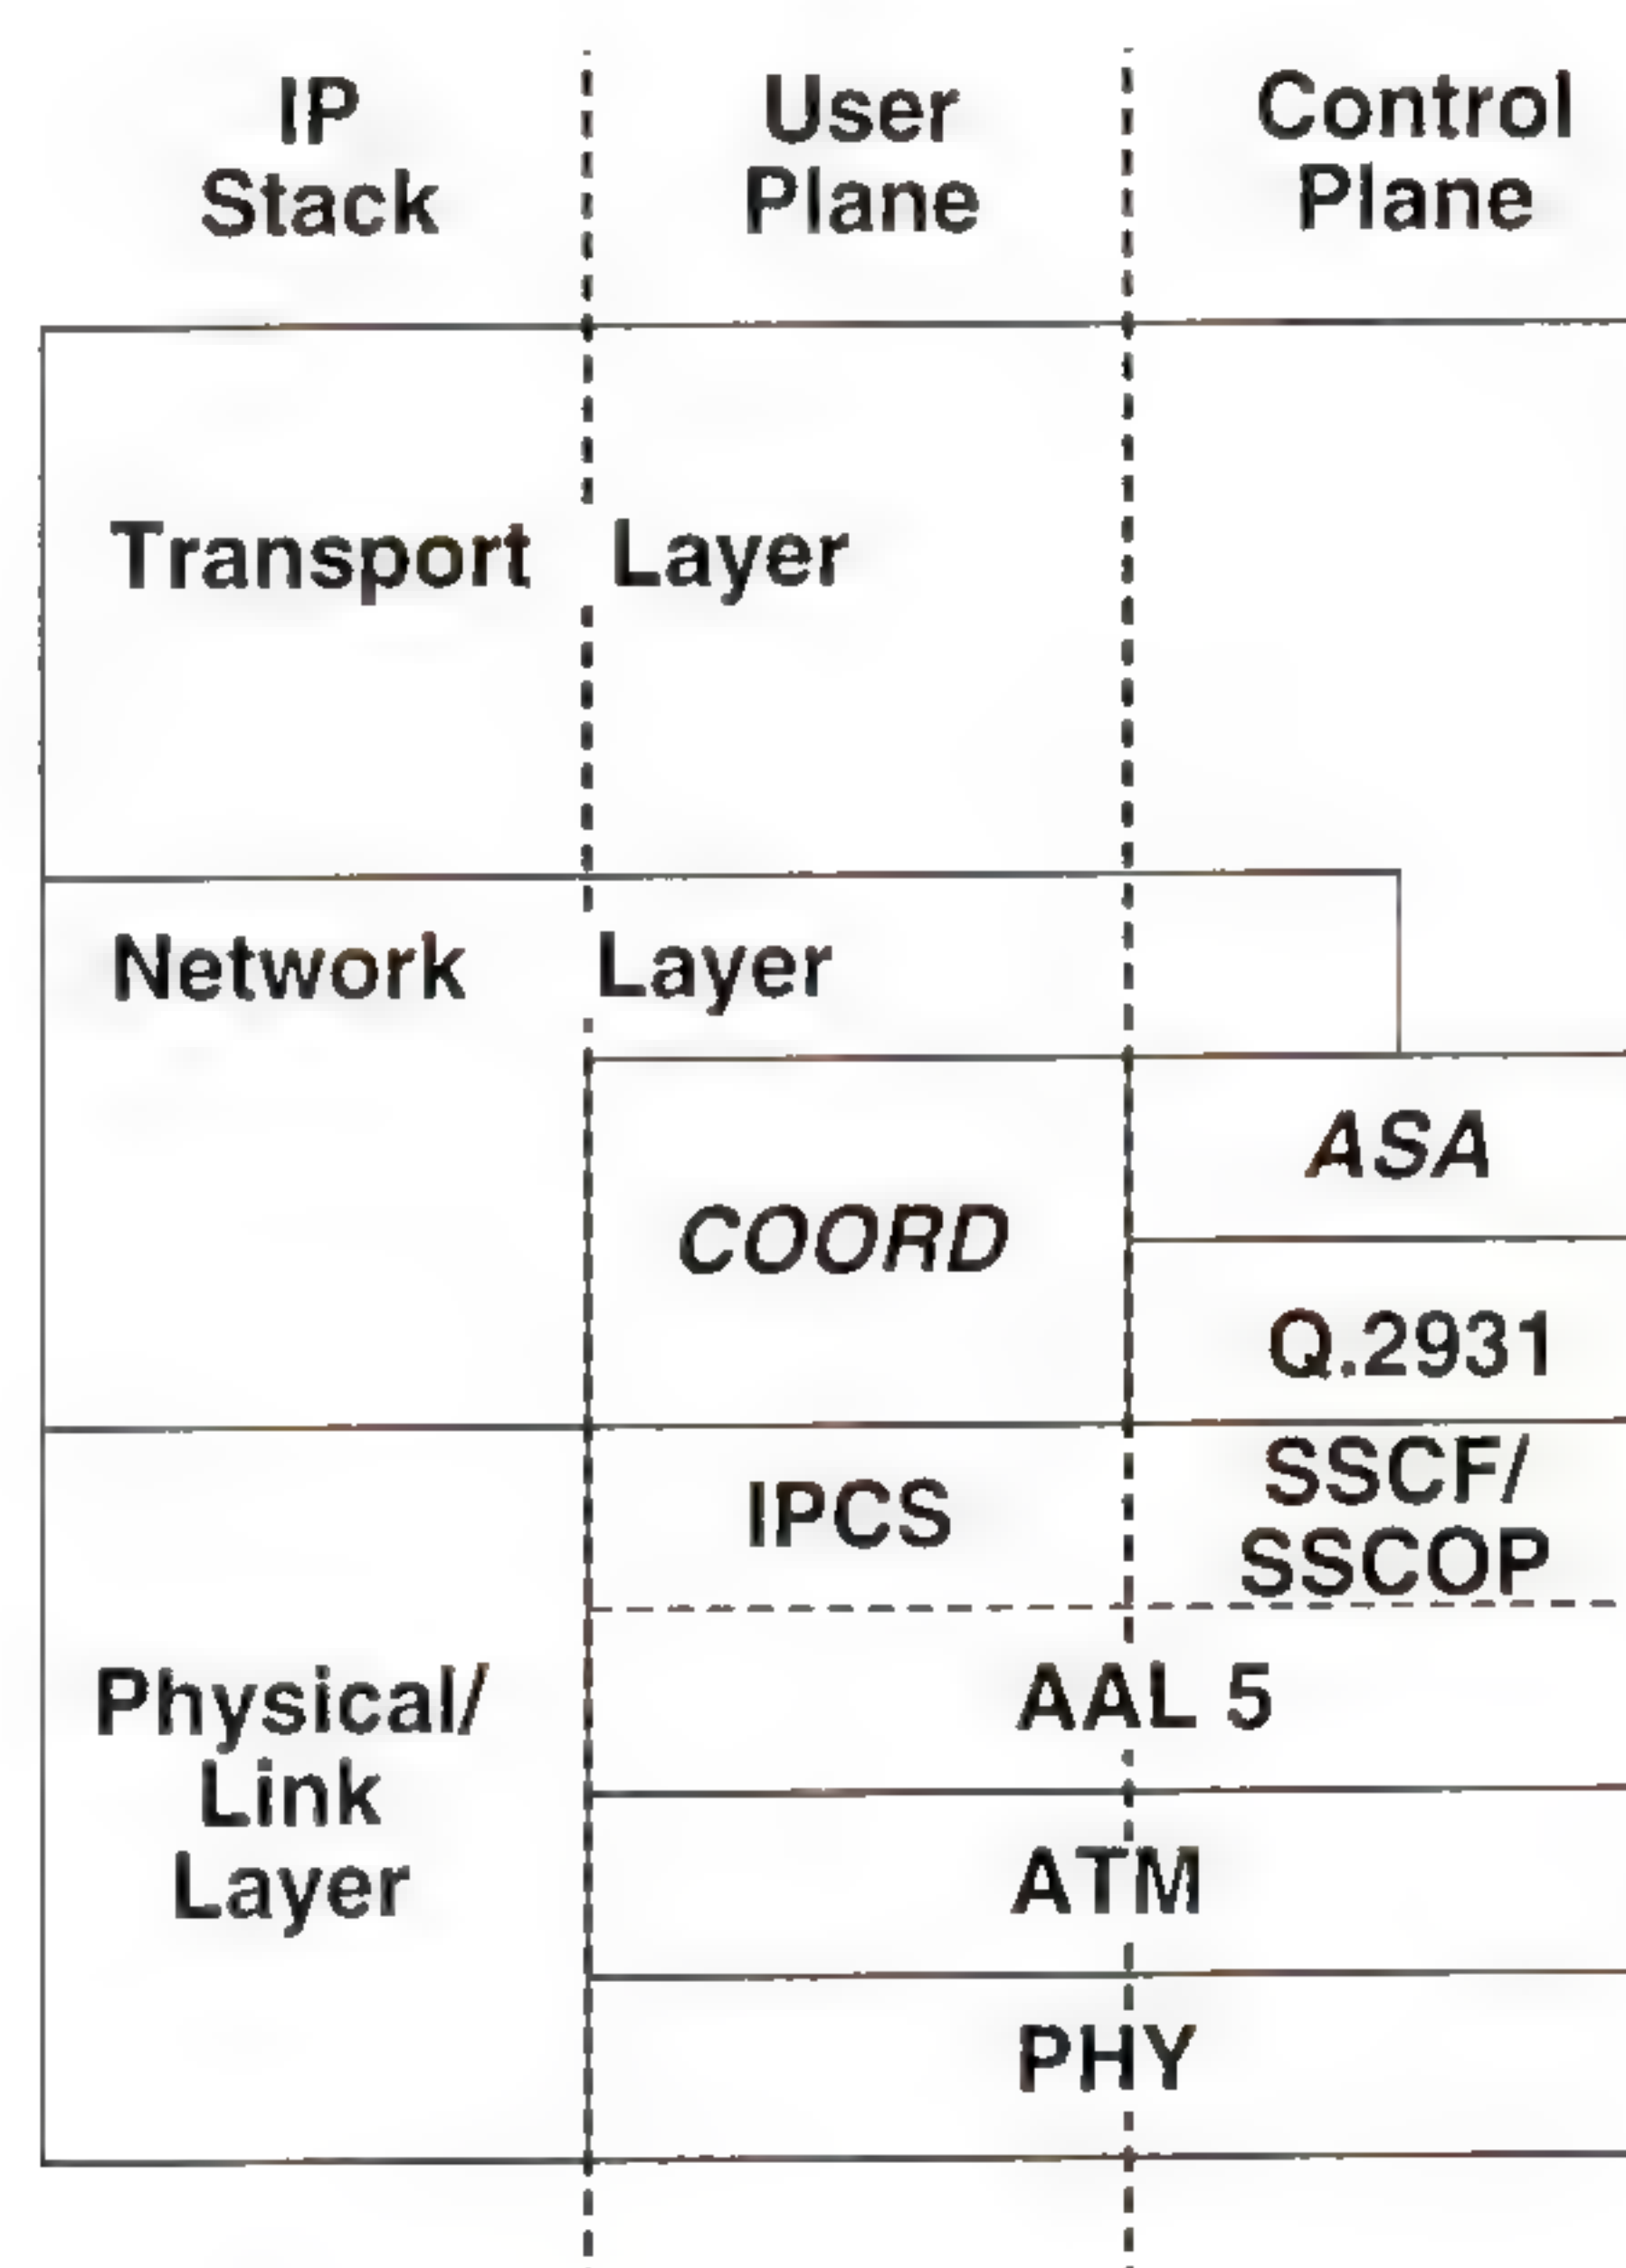| IP Stack	User Plane	Control Plane
Transport Layer		
Network Layer		
	COORD	ASA
		Q.2931
	IPCS	SSCF/ SSCOP
Physical/ Link Layer	AAL 5	
	ATM	
	PHY	

Figure 4. Protocol Stacks for Automatic Signalling Activation
(management Plane not shown)

Signalling for Data Services requires an additional layer in the user plane which communicates with the control plane. This can be a simple coordination process (COORD) indicating that a packet has arrived, and interrogating the control plane whether a connection already exists for the destination of the packet. If there is a connection, the packet can be forwarded immediately with the appropriate virtual channel identification. An IP specific *Convergence Sub-layer* (IPCS) performs the necessary encapsulation into AAL 5. If no connection is open, it has to be stored until a channel has been set up and sent afterwards. A generic protocol stack is shown in figure 4.

Signalling activation for UDP/IP

The ASA needs to know the IP destination address to determine whether a channel already exists to this address. This corresponds to a "called party number" in B-ISDN. As described earlier, the ASA has to have an address translation facility which could be used.

Interworking with B-ISDN *(continued)*

Maintenance of the address translation tables could be performed by an address resolution protocol or other management actions. In principle, the same connection can be used by different UDP/IP users once it is open, because connectionless services have no specific service requirements. The destination address is the minimum information needed for a basic operation. If additional services should be supported, like "address screening" of the IP addresses, the IP source address is needed. In addition to this, the Type of Service information of IP could be used in combination with a pre-configured database to determine the ATM QoS. The database would be needed to translate e.g., IP priority information like maximum throughput into precise ATM parameters like a corresponding peak bandwidth. Bandwidth allocation could also be based on well-known port numbers and should consider maximum bounds given by the link rate and the number of active connections on a link. Because there is no connection set-up nor release in UDP/IP, the connections have to be aged. Compared with the connection-oriented ASA, connection aging is the only additional functionality which is needed to support data services in B-ISDN. The described approach is similar to IP tunnelling in X.25 or N-ISDN [8], but it uses a more flexible bandwidth management.

Signalling activation for TCP/IP

For TCP/IP connections, the three-way handshake for connection set-up and release can be intercepted by ASA, again an additional functionality when compared with the connection-oriented case. This makes it unnecessary to age the connections. Apart from this, the same procedures are used as for UDP/IP signalling activation.

Alternatives for data services

Automatic Signalling Activation seems to be a possible candidate to challenge SMDS/CBDS services, as soon as signalling is introduced in B-ISDN. The reason for this observation is the minor changes required for a connection-oriented implementation. SMDS/CBDS has the advantage of faster routing decision after packet arrival. If the indirect support for connectionless services is used, the network resources are wasted such that this solution is only practicable for relatively small configurations. With the direct support of connectionless services, the number of permanent connections is lower, but network resources are still wasted when there is no significant traffic over the according links. Furthermore, it is very difficult for a public network provider to scale the bandwidth of such connectionless links properly. Another disadvantage of the SMDS/CBDS approach is that it is currently the only application to use AAL type 3/4.

Automatic Signalling Activation does not require much additional software in B-ISDN access equipment, once signalling has been introduced. The number of open connections over such an equipment might be restricted by the latency of signalling processing. An advantage of the ASA approach is that the same AAL type 5 can be used in the user plane and in the control plane.

Conclusion

It has been shown how B-ISDN signalling could support data services. As there are only minor changes compared to the implementation of connection-oriented services, it can be assumed that such implementations will challenge the current SMDS/CBDS service. It has been shown that the assumptions behind such an SMDS/CBDS service are imprecise because:

- Connection-oriented data services exist; the arguments for connectionless servers are based only on lower layer arguments
- Connectionless applications do not suddenly dump data at high speeds.

From a service viewpoint, connectionless servers might be superfluous. In a final comparison, for connectionless data services it has been shown that automatic signalling activation is a promising alternative to semi-permanent connections. Automatic Signalling Activation is therefore an important link for the full interoperability of existing data services with B-ISDN. An open question is still whether it enables all the supplementary services provided by ISDN.

References

[1] E. Biagioni et al. "Designing a Practical ATM LAN," *IEEE Network,* March 1993, pp. 32–39.

[2] Bellcore Special Report SR-NWT-002076, "Report on the Broadband ISDN Protocols for Providing SMDS and Exchange Access SMDS," Issue 1, September 1991, and related Technical References on SMDS: TR 772-775

[3] European Telecommunications Standards Institute, prETS 300217.1-4, "Connectionless Broadband Data Service," Parts 1 to 4, V27/11.92.

[4] ITU-TS, "Draft Recommendation I.364, Support of Broadband Connectionless Data Service on B-ISDN," COM 13-17-E, January 1994.

[5] ITU-TS, "Meeting Report—Services/Capability Sets/Workplan Group (12/6/93)," TD PL/11-13, Geneva, 29 November–17 December 1993.

[6] ITU-TS, "Q.2931 Overview," TD PL/11-55C, Geneva, 29 November–17 December 1993.

[7] The ATM Forum, "ATM User-Network Interface Specification, Version 3.0," Section 5, September 1993.

[8] Malis, A. G., Robinson, D., and Ullmann, R. L., "Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode," RFC 1356, August 1992.

[9] Laubach, Mark, "ATM for your internet—But When?" *ConneXions,* Volume 7, No. 9, September, 1993.

[10] ATM Forum, "User Network Interface (UNI) Specification Version 3.0," ISBN 0-13-225863-3, Prentice Hall, December 1993.

[11] Heinänen, Juha, "Multiprotocol Encapsulation over ATM," RFC 1483, March, 1993.

[12] Laubach, Mark, "Classical IP and ARP over ATM," RFC 1577, January, 1994.

[13] Atkinson, Ran, "Towards Real ATM Interoperability," *ConneXions,* Volume 7, No. 8, August 1993.

[14] Laubach, Mark, "IP Over ATM and the Construction of High-Speed Subnet Backbones," *ConneXions,* Volume 8, No. 7, July, 1994.

[15] Falk, G., "LAN Interconnection of X.25," *ConneXions,* Volume 7, No. 2, February 1993.

RETO BEELER received the Diploma in Electrical and Electronics Engineering from the Swiss Federal Institute of Technology (ETH) in Zurich, Switzerland in 1988. He joined Ascom in the same year, where he studied access protocols for fibre optic networks. In 1991 he was an academic visitor to the Computer Laboratory of the University of Cambridge (UK). He currently works in the Broadband-ISDN department of Ascom corporate research on interworking aspects between MANs and ATM networks. E-mail: `beeler@tech.ascom.ch`

The INRIA Videoconferencing System (IVS)

by Thierry Turletti, INRIA

Introduction

The *INRIA Videoconferencing System* (IVS) is a software system to transmit audio and video data over the Internet. It includes PCM and ADPCM audio codecs, as well as a H.261 [5] codec. Both audio and video codecs are software codecs.

The H.261 video coding standard was originally designed for the transmission of video flows over fixed-bandwidth lines, i.e., leased lines or switched circuits for data transmission. In order to use this video coding over packet switched networks such as the Internet, a packetization scheme is required. This scheme must take into account the hierarchical structure of H.261 images. Furthermore, packet loss recovery and flow control schemes are also required to send video over the Internet. We next briefly describe the H.261 coding standard. We then present the video packetization scheme and the error and flow control schemes we developed for the Internet environment. We then describe different applications using IVS. We conclude with practical details regarding IVS, in particular the platforms supported by the current version, where to get it, etc.

H.261 video coding

The H.261 recommendation includes descriptions of a coding mechanism and a scheme to organize video data in a hierarchical fashion. The H.261 video coding uses state of the art video compression encoding methods. The high compression rate obtained by the coding mechanism allows running videoconferencing at very low data rates. The compression techniques include a *Discrete Cosine Transform* (DCT), a quantization, Huffman encoding, and optionally vector motion compensation. The structure of the coding system is shown in Figure 1.
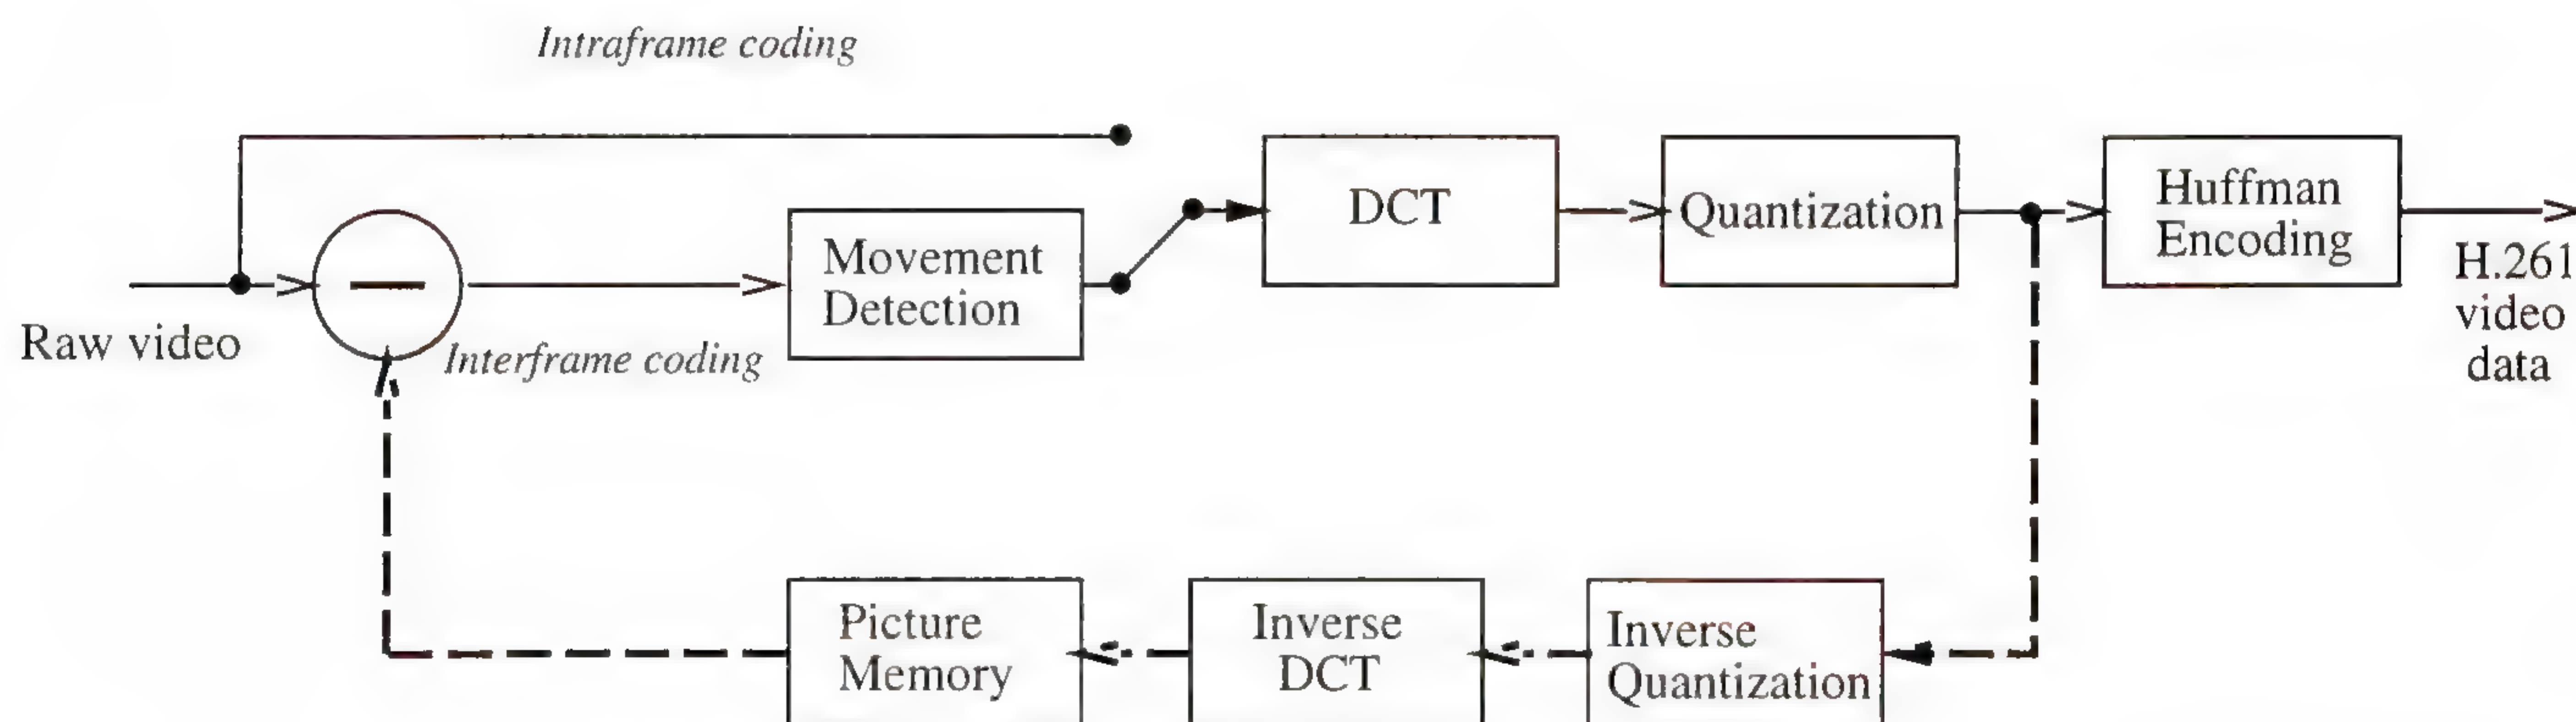

Figure 1: H.261 Video Coding System

The *International Telecommunication Union* (ITU) has adopted standard picture formats including the *Common Intermediate Format* (CIF) and the *Quarter-CIF* (QCIF). CIF has 352 pixels per line and 288 lines per picture. QCIF has half as many pixels and half as many lines as CIF. H.261 coding is done either on the input pictures, or on the difference between successive images. The first case is referred as *intraframe* coding, the second case as *interframe* coding. Intraframe coding means that the image is encoded without any relation to the previous images. This kind of encoding removes only the spatial redundancy in a picture, whereas interframe coding also removes the temporal redundancy between successive pictures. In the latter case, the difference between the current and the predicted image is DCT transformed and then linearly quantized.

CIF and QCIF pictures are arranged according to a hierarchical structure that includes four layers, namely the *Picture* layer, the *Group Of Blocks* (GOB) layer, the *MacroBlock* (MB) layer and the *Block* (B) layer (see Figure 2).
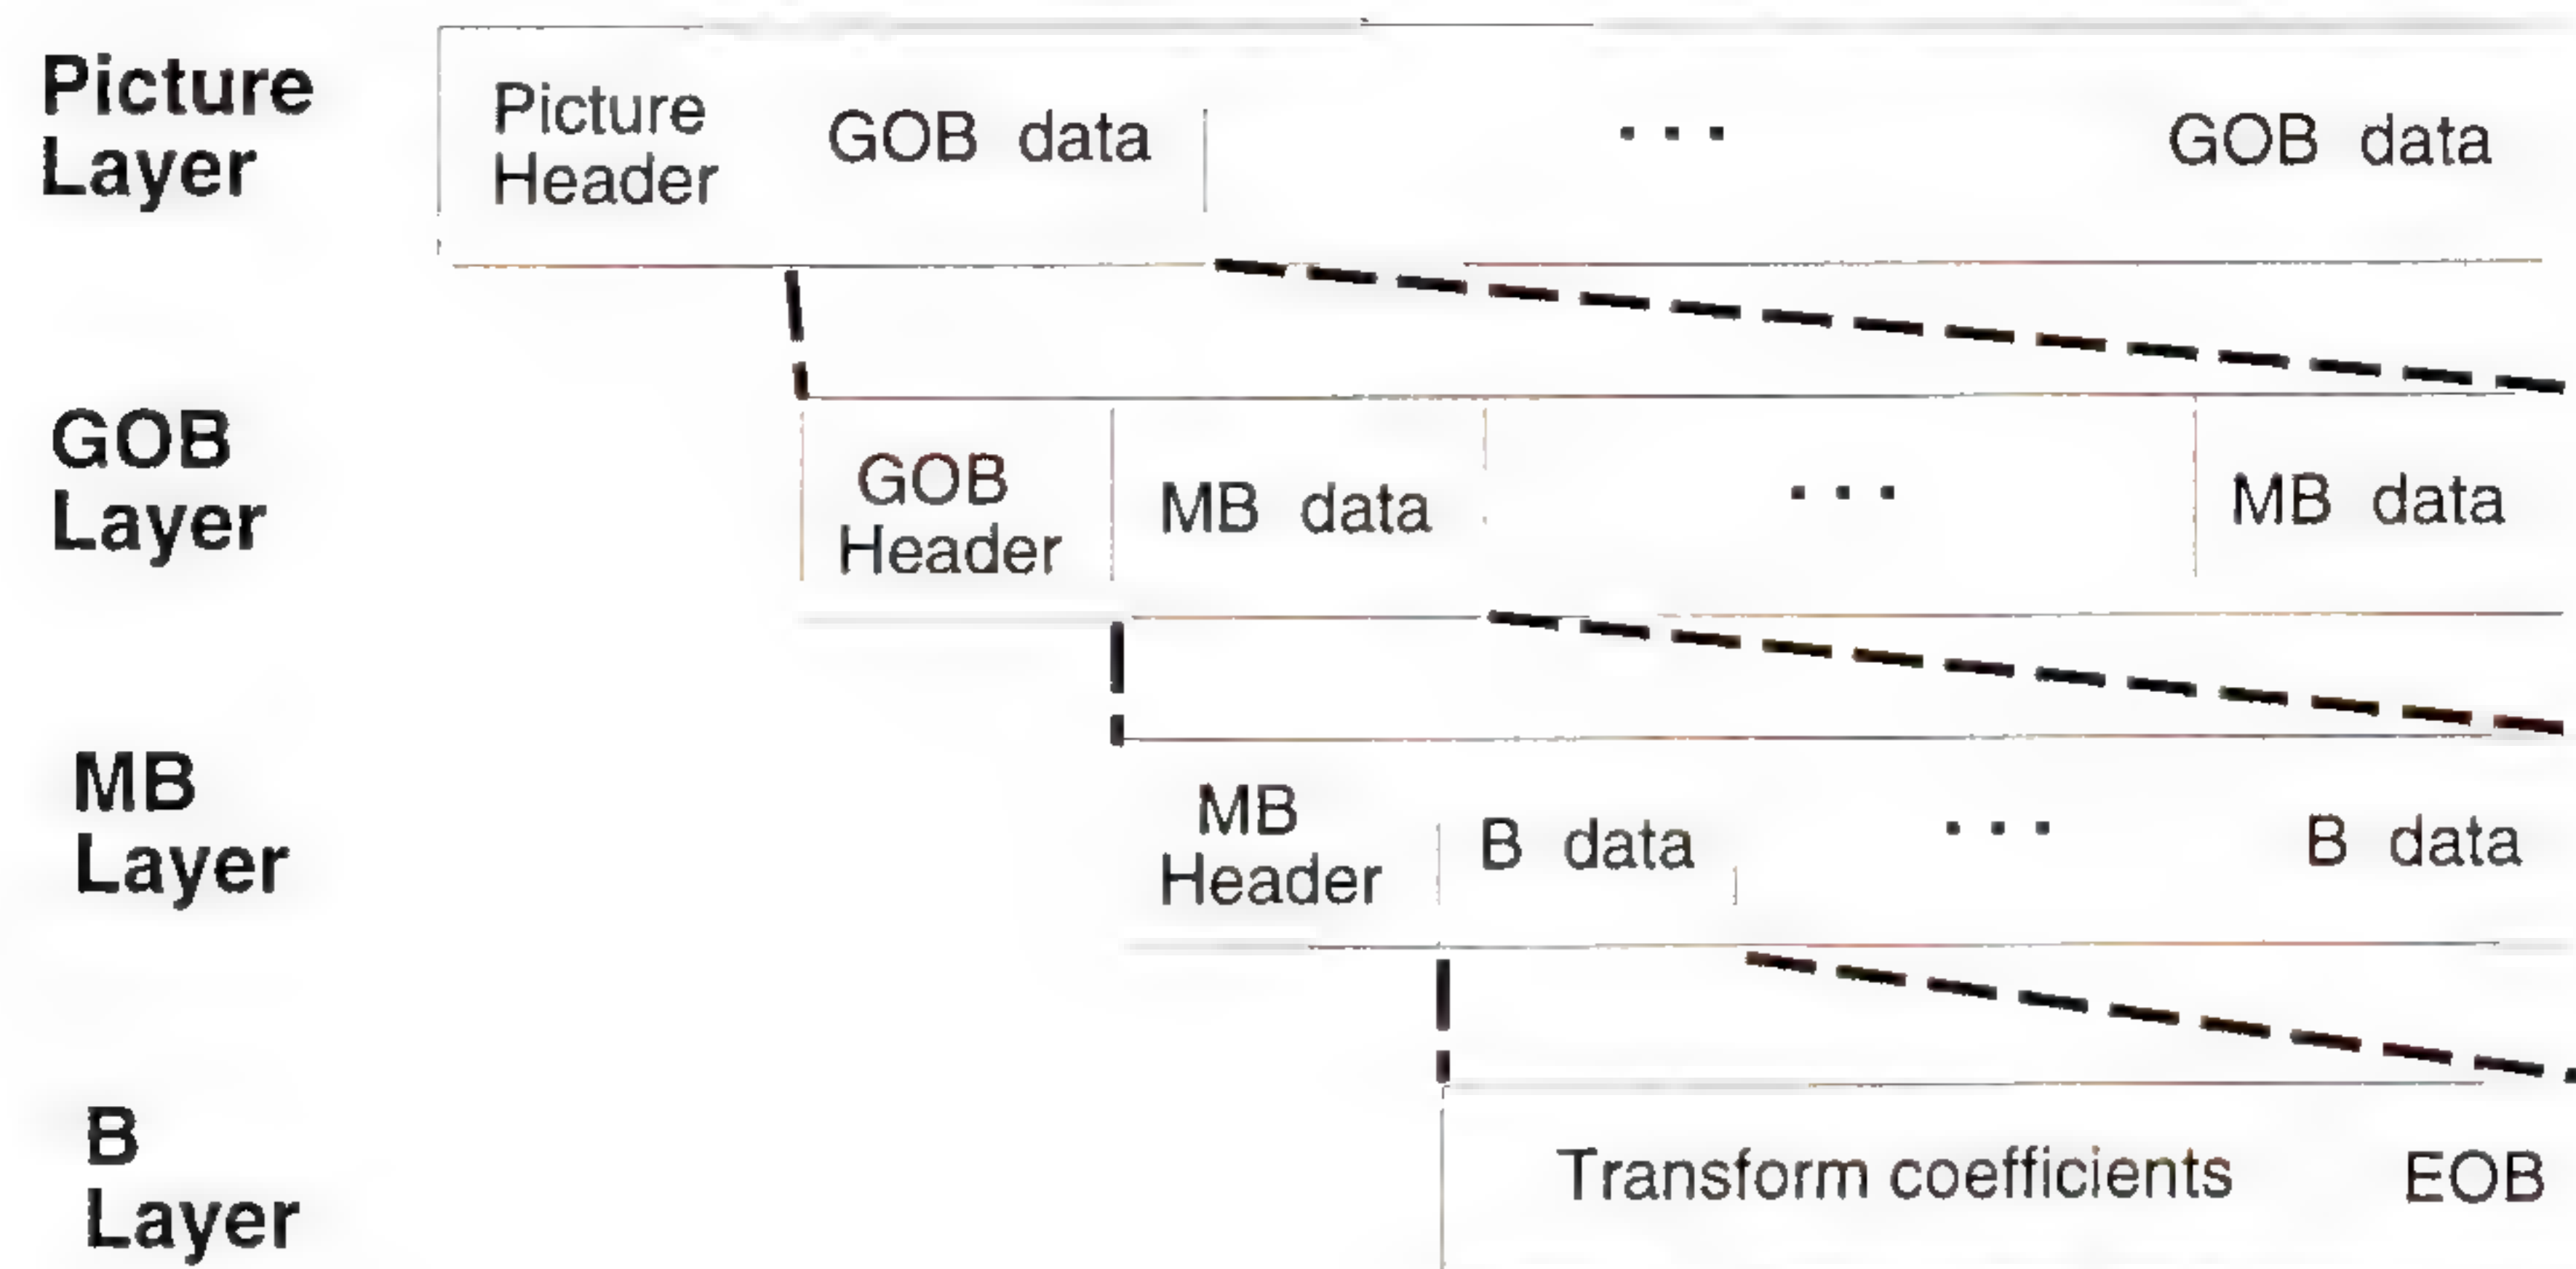

Figure 2: CIF and QCIF layers

A CIF picture is divided into 12 GOBs, while a QCIF picture is divided into 3 GOBs. A GOB is composed of 3 rows of 11 MBs and each MB contains six blocks of 8x8 coefficients. Data for a Block consists of codewords for transform coefficients followed by a fixed-length code End of Block (EOB) marker. All of the quantized coefficients are ordered into a "zigzag" sequence which helps to facilitate entropy coding by placing low frequency coefficients (which are more likely to be non-zero) before high frequency coefficients.

The transmitted bitstream contains a BCH (511, 493) forward error correction code (FEC), but its use by the decoder is optional.

Sending H.261 video over the Internet

To send H.261 video using standard UDP datagrams, a packetization scheme for H.261 video streams has been described in an Internet Draft [9] submitted to the *Audio Video Transport* (AVT) Working Group at the IETF. RTP [4] is used over UDP to achieve multiplexing. The proposed scheme takes specifically into account the hierarchical structure of the H.261 coding. The bitstream produced by a standard H.261 coder includes forward error correction (FEC). The FEC is over 492 bit blocks of the encoded bit stream, and it bears no relation to the hierarchical structure of H.261, i.e., Picture, GOB, MB layers. So, to conform to the Application Layer Framing (ALF) philosophy, and in order to break the bitstream into units which do not have any dependencies on other parts of the bitstream, the FEC is removed. The smallest unit of data that has the fewer dependencies on other parts of the bitstream (ADU) is the GOB. Indeed, MBs have addresses relative to preceding MBs, and quantizer and motion compensation vector of a MB may depend on a previous MB within the same GOB.

The output data flow generated by a H.261 coder is a variable rate flow. The intensity of the flow depends on the quality of the video camera, the type of the images being encoded which is a function of the movement, the scene structure, the scene lighting, etc. Most of the time, GOB information fits in a single packet. However, if scenes changes are few, several GOBs can be grouped inside the same packet to decrease the packet sending rate and thereby avoid network congestion. A small header is added to each packet to encode information required to decode out-of-order ADUs received and to efficiently depacketize segmented GOBs received.

21

INRIA Videoconferencing System *(continued)*

As part of the MICE (*Multimedia Integrated Conferencing for Europe*) project [10], this packetization scheme has been used by some commercial hardware codecs (in particular GPT, and Bitfield codecs) allowing interoperability between these codecs and IVS. (See URL <http://www.cs.ucl.ac.uk/mice/mice.html>).

IVS also includes an error control scheme to deal with lost ADUs. The interframe coding is much more sensitive to packet loss than the intraframe coding. The visual error due to a lost ADU will persist from frame to frame until the corrupted portion of the image is refreshed in intraframe coding. The fastest way to correct the error is to request the coder to force intraframe coding of the part of the image blurred by the loss. Note that this is not a retransmission but only a refreshment since the encoding occurs for a new frame. The request to refresh (in practice a NACK, i.e., negative acknowledgment) should be sent as soon as the loss is detected to speed up the error recovery. However, to prevent the NACK implosion problem, this method is only efficient when the number of receivers is small. In IVS, NACK packets are used only if there are less than 10 participants in the conference. When NACKs cannot be used, another solution is to increase the intraframe coding refreshment rate. In IVS this rate currently depends to the loss rate observed in the network. If the coder is not able to adjust its intraframe encoding rate (as in the case of a hardware codec), it can either use intraframe-only coding or periodically request a full intraframe encoded image.

In order to prevent swamping all the resources of the Internet, IVS includes a feedback control mechanism [1] in which the parameters of the coder are adjusted according to the network conditions observed. The output rate control is adjusted by changing either the video frame rate or the quantizer value and the movement detection threshold. The specific requirements of a video application will indicate which of the three parameters should be modified when adjusting the output rate of the coder. The frame rate is modified if the precise rendition of individual images is important. The quantizer and the movement detection threshold are changed if the perception of movement is more important. The feedback information consists in the loss rate indication sent back from decoders to the coder. The feedback mechanism [2] uses a novel probing mechanism to solicit feedback information in a scalable manner.

IVS applications

Videoconferencing can be used to support many applications such as holding working meetings between distant sites, remote teaching and seminars or conference broadcasts. For example, IVS is currently used by the MICE partners located in Belgium, France, Germany, UK, Sweden and Norway, both for weekly meetings and MICE seminars [7]. SD [6] is often used to advertise a session which is to be multicast. It provides users with a convenient way to join a session simply by selecting the session entry from a list of advertised conferences.

You can also use IVS to call up someone on the network, provided that the correspondent is running an IVS daemon (*ivsd*). You only have to start a unicast IVS with the correspondent's machine address and click on the "Call up" button. The distant daemon will ring and pop up a message showing your talk request. If the correspondent clicks on the "Accept" button, a unicast IVS session will be automatically started and the videoconference can begin.

If the correspondent is not present, a message will remain on the correspondent's screen showing your name and the time of the request. It is also possible to record and replay a broadcast conference or an audio/video clip. The clip obtained can also be sent by e-mail or be played back via *Mosaic* using *Metamail* MIME extensions.
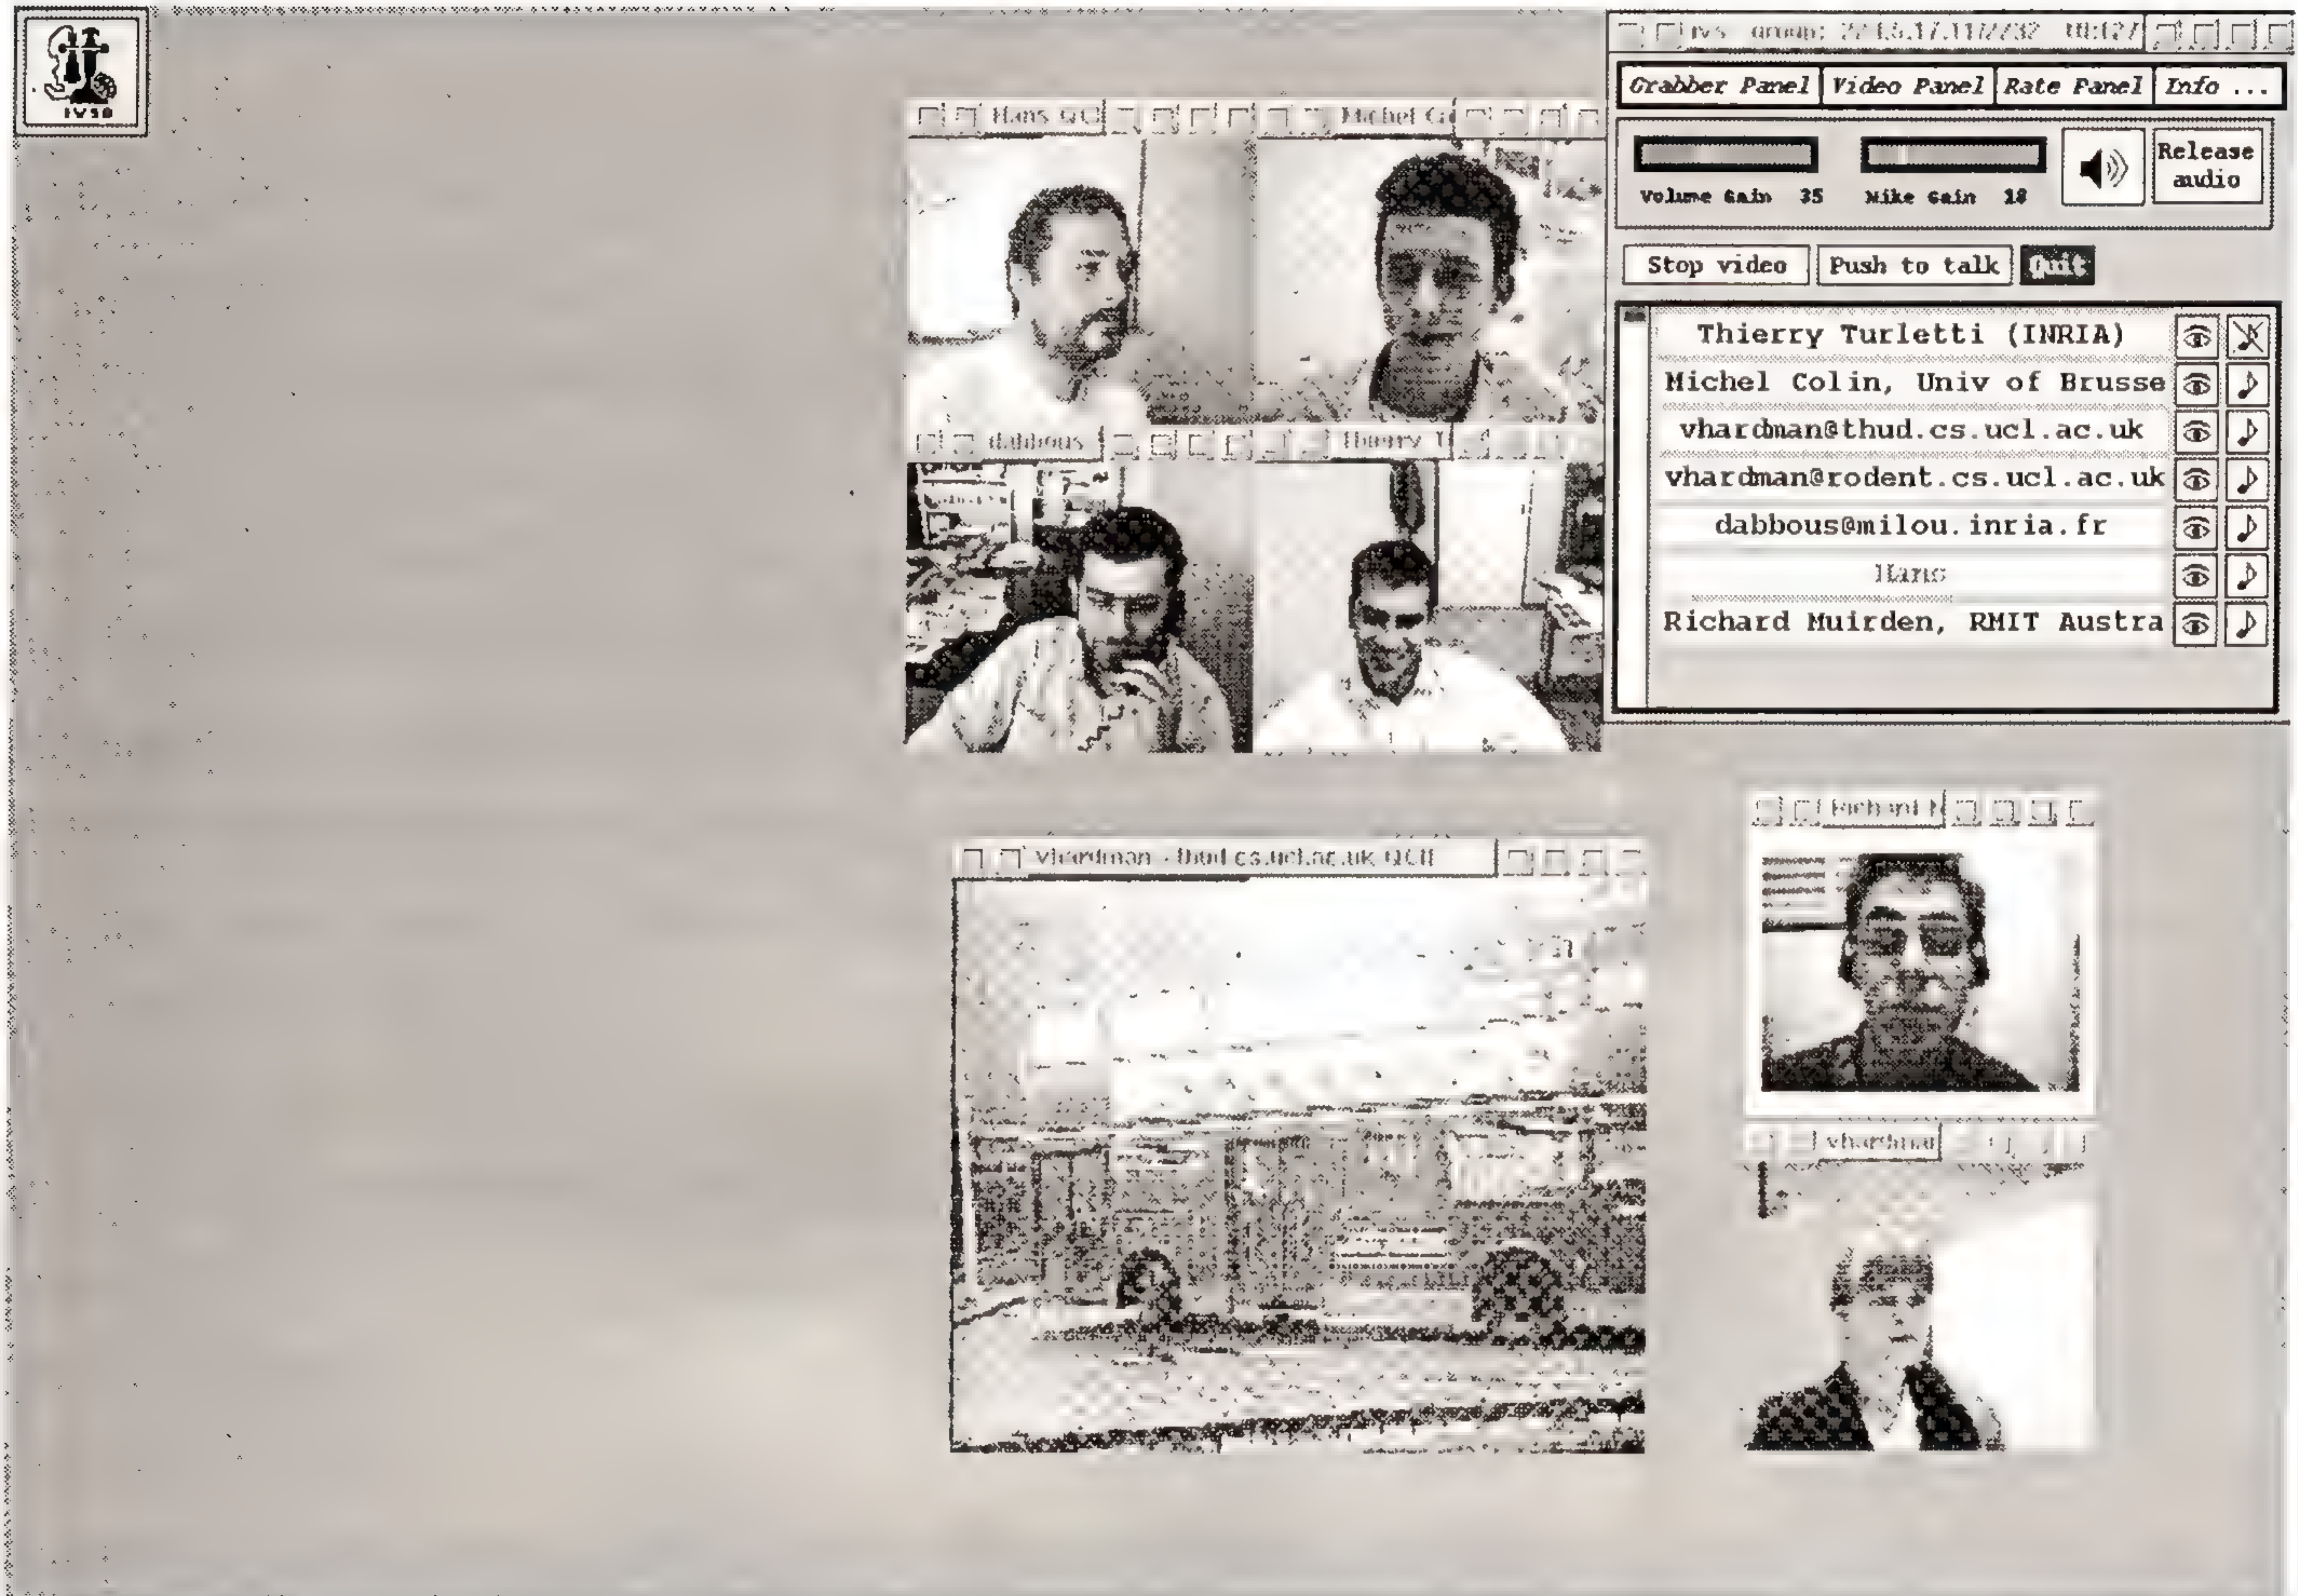

Figure 3: Screendump from a typical IVS session

Platforms IVS currently runs on the following platforms:

- SUN workstations under SunOS 4.1.3 or SunOS 5.3 with VideoPix, Parallax and SunVideo grabbers.

- SGI IRIS Indigo and Indy workstations under Irix 5.2 with Galileo and Indigo grabbers.

- DEC-5000 workstations under Ultrix 4.3 with VideoTX grabber and Amaxine or Alofi audio servers.

- HP9000/7xx workstations under HP-UX 8.0x with HP VideoLive grabber.

IVS requires a workstation with a screen with 1, 4, 8 or 24 bits depth. Multi-host conferences require that the kernel support multicast IP extensions [3]. Sources and binaries are available by anonymous ftp from: `zenon.inria.fr in rodeo/IVS/last_version`.

Future work We are currently studying the use of packet-level forward error correction (FEC) techniques to enhance audio quality in cases of packet loss.

23

INRIA Videoconferencing System *(continued)*

We are also looking at the possibility of adapting the audio compression ratio to variable network transmission capacities. Work is underway to develop versions for PC/Windows and DEC/J300 platforms.

References

[1] J-C. Bolot, T. Turletti, "A rate control for packet video in the Internet," Proc. IEEE INFOCOM '94.

[2] J-C. Bolot, T. Turletti, I. Wakeman, "Scalable feedback control for multicast video distribution in the Internet," Proc. SIGCOMM '94.

[3] S. Deering, "Host extensions for IP multicasting," RFC 1112, August 1989.

[4] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A transport protocol for real-time applications," Internet-Draft, July 1994.

[5] Recommendation H.261: "Video codec for audiovisual services at px64 kbits/s," International Telecommunication Union, 1993.

[6] V. Jacobson, "SD manual pages," Lawrence Berkeley Laboratory (LBL), March 1993.

[7] A. Sasse, U. Bilting, M. Handley, C.D. Schulz, T. Turletti, "Remote Seminars through Multimedia Conferencing: Experiences from the MICE project," Proc. of INET '94/JENC5, 1994.

[8] T. Turletti, "A H.261 software codec for videoconferencing over the Internet," INRIA Research Report, No. 1834, Jan. 1993.

[9] T. Turletti, C. Huitema, "Packetization of H.261 video streams," Internet-Draft, September 1994.

[10] Crowcroft, J., "MICE: Multimedia Integrated Communication for Europe," *ConneXions,* Volume 7, No. 11, November 1993.

[11] S. Deering, "IP Multicasting," *ConneXions,* Volume 5, No. 3, March 1991.

[12] R. G. Herrtwich and L. Delgrossi, "The ST-II Protocol for Multimedia Communication," *ConneXions,* Volume 8, No. 1, January 1994.

[13] M. Macedonia and D. Brutzman, "MBone Provides Audio and Video Across the Internet," *IEEE Computer,* April 1994.

[14] D. Clark, "The Design Philosophy of the DARPA Internet Protocols," in the Proceedings of ACM SIGCOMM '88.

[15] S. Casner, and S. Deering, "First IETF Internet Audiocast," *ConneXions,* Volume 6, No. 6, June 1992.

[16] Borenstein, Nathaniel S., "*Metamail*—Multimedia Mail for the Masses," *ConneXions,* Volume 6, No. 3, March 1992.

[17] Vaudreuil, G., "MIME: Multi-Media, Multi-Lingual Extensions for RFC 822 Based Electronic Mail, *ConneXions,* Volume 6, No. 9, September 1992.

[17] M. J. Handley, P. T. Kirstein, M. A. Sasse, "Multimedia Integrated Conferencing for European Researchers (MICE): piloting activities and the Conference Management and Multiplexing Centre," *Computer Networks and ISDN Systems,* No. 26, 1993.

THIERRY TURLETTI received an M.Sc. in Computer Science and Digital Signal Processing from University of Nice-Sophia Antipolis in 1990. Since November 1991, he has been a Ph.D student at INRIA Sophia Antipolis specializing in multimedia systems. He has been a member of the MICE project since December 1992. His areas of interest include audio and video compression and control congestion algorithms. He can be reached at: `turletti@sophia.inria.fr`

Summer IETF Meeting Security Report

by Jim Galvin, Trusted Information Systems

At the Toronto *Internet Engineering Task Force* (IETF) meeting (July 25–29), some new developments were announced, including the formation of a Directorate, the adoption of a proposal to create a working group to address the Internet's key management needs, and the drafting of a document to express the "vision" for the security area.

Formation of a Directorate

Forming a Directorate is a step undertaken in many IETF areas. The members of the Directorate are security experts who will assist the Security Area Director (currently Jeff Schiller of MIT) in reviewing specifications and offering security advice to IETF working groups. The members of the security directorate are:

Jeff Schiller	`<jis@mit.edu>`
Jim Galvin	`<galvin@tis.com>`
Ran Atkinson	`<atkinson@itd.nrl.navy.mil>`
Steve Bellovin	`<smb@research.att.com>`
Steve Crocker	`<crocker@tis.com>`
Barbara Fraser	`<byf@cert.org>`
Phil Karn	`<karn@qualcomm.com>`
Steve Kent	`<kent@bbn.com>`
John Linn	`<linn@ov.com>`
Cliff Neuman	`<bcn@isi.edu>`
Rob Shirey	`<shirey@mitre.org>`
Ted Ts'o	`<tytso@mit.edu>`

Directorate meetings will be closed; the Security Area Advisory Group (SAAG) meetings, held late in the week of the week-long IETF meetings, will continue to provide an open forum for reporting on the security area and introducing new ideas.

A "Vision" for the Security Area

The Security Area Vision will be drafted by Schiller, reviewed by the SAAG, and, when adopted, published as an informational document. It will help guide the strategic activities of the Security Area.

IPNG Activities

The *Internet Protocol Next Generation* (IPNG) Area announced its strategic direction for the next generation of IP, a critical element of which is that security be available by default. Three cryptographically secure services are envisioned: data integrity, data origin authentication, and data confidentiality.

Key Management Working Group

Adding security services to existing protocols and including them in new protocols share a common need: key management. Many of the mechanisms that provide authentication and encryption services require a mechanism to distribute either a secret value, a cryptographic key, or both. In response to this common need, the security area will create a key management working group. One of the options the group will consider will be the use of the Domain Name System to distribute public keys when asymmetric cryptography is used.

DNS Security Working Group

The *Domain Name System* (DNS) Security working group met to discuss the two proposals currently under evaluation for adding digital signature services to the DNS protocol. The group will continue its evaluation and choose a proposal on the DNS Security mailing list, after which TIS will begin implementing the chosen proposal.

More Information

For information about DNS security join the `dns-security@tis.com` e-mail list by sending a subscription request e-mail message to: `dns-security-request@tis.com`. For information about the security area join the `saag@tis.com` e-mail list by sending a subscription request message to: `saag-request@tis.com`.

Call for Papers

The *6th Joint European Networking Conference* (JENC6) will be held in Tel Aviv, Israel, May 15–18, 1995. The conference will address strategic aspects of open computer networking. It aims to bring together people involved in computer networking for research and education, from all parts of the field: industry, government, research, and education. The conference is being organised by the new Association resulting from the merger of RARE and EARN (*Réseaux Associés pour la Recherche Européenne,* and *European Academic & Research Network*), and will take place in the new Dan Panorama Convention Center in Tel Aviv, Israel.

Background

European research networking has moved far beyond its early deployment in academic environments and has left past controversies such as low-level protocol choice far behind. It is now in a period of continuing development of infrastructure, expansion of the initial base of research users into broader communities, and greater use of new and better network services. The goal of this conference is to survey the current situation, illuminate major unresolved issues and technologies, and to engage participants in discussion on possible future directions.

As in previous years, the conference will address the requirements of national and international service providers of different kinds, as well as application developers and user groups of various disciplines. Information Technology experts interested in the broad impact of changes in current and upcoming technologies, and professionals working in the area of research network funding and management, will have a unique opportunity to exchange views and meet with the key players in this fascinating, continuously changing world.

Theme

The conference theme, "Bringing the World to the Desktop," is our metaphor for two major changes under way: the first is the increasing penetration of daily research/educational work and practices by networks and networking technology; the second is the new set of requirements that desktop networking implies for the underlying technology and the structures of service provision. Papers which reflect this theme will be favoured by the selection committee. The official language of the conference will be English.

Topics

Topics for presentations include, but are not limited to, the following:

- *Networking Technology and Engineering:*
  High Speed Networking
  Desktop/WAN integration
  Routers/Switches
  Home/Mobile Networking
  IPng Migration

- *Support for Cooperative Work:*
  Compound Document/Information Architectures
  Teleconferencing
  Group Calendaring and Scheduling Systems
  Shared Editing and Shared Whiteboards
  Workflow Applications
  Support for Asynchronous Interaction
  Inter-Application Communication

- *Security and Privacy:*
  Authentication and Integrity
  Privacy and Confidentiality
  Anonymity
  Reliable Telecooperation
  Business over Open Networks
  Secure/Insecure Applications
  Management of Security
  Policy Issues
  Firewalls
  Cryptographic Mechanisms (e.g., digital signature/encryption)
  Physical Security

- *Providing and Accessing Information:*
  Tools (*Gopher, WAIS, WWW,* etc.)
  Proxies, Caching and Mirroring
  Online services, (e.g., Electronic Publishing, Libraries, Databases)
  Directory Information
  User Support

- *Policy Related Issues:*
  Regional Issues
  Funding Models
  Global Top-level Interconnect
  Social Implications
  Influence of Governments

Submissions

All papers must be written in English and should be between 2000 and 4000 words. Each paper should include a 200 word abstract, 3 to 5 keywords, and the full name(s) and address(es) of the author(s).

Electronic submission is highly recommended. Plain text or *Post-Script* files can be sent by e-mail to: `jenc6-submit@rare.nl`. *Post-Script* files can be sent by anonymous FTP to `erasmus.rare.nl` (IP address 192.87.30.2), into directory `pub/jenc6/submit`. Please note that files deposited in this directory can only be written once and cannot be deleted afterwards. Should electronic submission be impossible, please submit 6 copies of a full paper manuscript to the address below.

Important dates

Full manuscripts due:	November 25, 1994
Proposals for demonstrations due:	December 30, 1994
Notification of acceptance to authors:	January 20, 1995
Camera-ready papers due:	March 31, 1995

More information

To be added to the conference e-mail distribution list send a message to: `jenc6-request@rare.nl`. For more information contact:

RARE Secretariat
Singel 466-468
NL-1017 AW Amsterdam
The Netherlands
Tel.: +31 20 639 1131
Fax: +31 20 639 3289
E-mail: `jenc6-sec@rare.nl`

Call for Papers

The 5th *International Workshop on Network and Operating System Support for Digital Audio and Video* (NOSSDAV '95) will be held April 19–21, 1995 in Durham, New Hampshire, USA. The event is sponsored by the IEEE Communications Society in cooperation with ACM SIGCOMM, SIGOPS, SIGMM, SIGGRAPH, and SIGIR.

Background

Network and operating system support for digital audio and video are becoming increasingly important with the convergence of the computer, the TV, and communications. Innovation in this field is fueling the industry developments in interactive multimedia services to the home. In the past, research in this domain has largely originated as adaptations of specific technologies to support audio and video. This work has lead to an understanding of common cross-disciplinary problems. Increasingly, this work encompasses and integrates the diverse technology necessary to achieve end-to-end systems. To this end, research leading to complete solutions is viewed as particularly important to the workshop.

Topics

The workshop is intended to bring together practitioners from a variety of areas, including communications and networks, operating systems, real-time systems and distributed computing. It is intended that an outcome of the workshop will be a statement of the the state of the art in this field, highlighting the major areas requiring future research. Relevant topics for the workshop include:

- High-speed/ATM networks
- Multimedia-oriented desk, local, and wide area networks
- Workstation architectures for multimedia
- Cell-based system architectures
- Multimedia network interfaces
- Communication protocols for multimedia
- Multicast
- Micro-kernel and OS support for real-time communications
- Resource management and reservation in the OS and network
- End-to-end admission control
- Quality of service and synchronization frameworks
- Multimedia storage, server, and I/O architectures
- Distributed multimedia systems
- APIs and CM programming abstractions for multimedia
- Community networking
- TV set-top device communication

Submissions

Two types of submissions are solicited: *position papers* and *research papers*. For the purpose of paper review, position papers are restricted to three single-spaced ASCII pages. Research papers are restricted to an extended abstract no longer than five formatted *PostScript* pages. Papers should be sent via e-mail to: `nossdav95@spiderman.bu.edu` Only if electronic submission is impossible, papers may be sent to:

Prof. T.D.C. Little
Multimedia Communications Laboratory
Department of Electrical, Computer and Systems Engineering
Boston University
44 Cummington Street
Boston, MA 02215
Tel: +1 617 353-9877 • Fax: +1 617 353-6440 • E-mail: `tdcl@bu.edu`

Important dates

Abstracts due:	December 12, 1994
Acceptance notification:	January 25, 1995
Final paper due:	March 8, 1995

Call for Papers

Interop Company is soliciting technical papers for an *Engineer's Conference* to be held as part of the upcoming *NetWorld+Interop 95* event, March 27–31 1995, in Las Vegas, Nevada. The Engineer's Conference, which will run March 29–30, is a two-day focused event offering approaches and solutions to practical systems and software design for networking. All conference participants will be able to attend the *NetWorld+Interop 95* exhibition, which runs March 28–30.

Format
The conference will feature the presentation of original papers which will have been selected by a review committee. All accepted papers will be published in Conference Proceedings. Accepted papers must be presented by original authors during the 2-day conference.

Topics
The Engineer's Conference will concentrate on engineering design problems in three areas: *High Speed Networking, Internetworking,* and *Network Management.* This conference seeks to bring together research scholars, engineers, and vendors to address pragmatic engineering issues in the field of networking and distributed systems interoperability. It is an excellent forum for engineers and researchers to publish papers on solutions to today's engineering-related problems. Papers are solicited in the following areas:

- *High Speed Networking:* ATM, Fast Ethernet, SONET, FDDI-II, HIPPI, SMDS, Frame Relay, Broadband ISDN, etc.

- *Internetworking:* Design of Bridges, Routers, and Multiprotocol Brouters, Addressing Schemes, Routing Protocols, Application Gateways etc.

- *Network Management:* Bandwidth utilization, Traffic Trend Analysis/Capacity Planning, Automated Trouble Ticket Systems, Congestion detection, Network Simulation, SNMP v1 and v2, Security, Export considerations for secure systems, Manager-to-manager communications, Standardized Testing Suites, Expert Systems, Accounting, Distributed/Hierarchical Management architectures, etc.

Submission Guidelines
Interested authors are invited to submit an abstract (up to 100 words) clearly describing the problem and the solution offered. All abstracts will be reviewed and authors will be notified for acceptance or rejection of the abstract. Authors of accepted abstracts must submit the paper before the last date. These papers are reviewed by a technical committee for technical merit of the paper before final acceptance. All abstracts must contain the authors name, address, telephone number, Fax number and e-mail address (if available). Please send your abstract to:

Interop Company
303, Vintage Park Drive Suite 201
Foster City, CA 94404–1138
USA
Attn.: Engineer's Conference

or e-mail it (in ASCII or *PostScript*) to: `engineer@interop.com`

Important dates

Abstracts due:	October 1, 1994
Notification to Authors:	October 15, 1994
Draft paper due:	December 1, 1994
Feedback to authors:	December 24, 1994
Camera ready copy due:	January 10, 1995
Overhead slides due:	February 15, 1995

NetWorld+Interop 95 moves from Berlin to Frankfurt

As from 1995, *NetWorld+Interop* will take place in Frankfurt and no longer in Berlin. Next year, this event will be held from May 29th through June 2nd.

Many things contributed to the decision of the organizers Ziff Messe & Konferenz to move from the exhibition site in Berlin where last June the European premiere of NetWorld+Interop was held with great success.

More than 20,000 professionals and over 1,100 tutorial and conference participants visited the exhibition. The 311 exhibitors and the strategic partners Ziff Messe & Konferenz and Novell were exceptionally happy with the excellent results achieved at this premiere.

But despite this success in Berlin, the planned venue for the next NetWorld+Interop, which was supposed to be held between the 28th and 30th of June 1995, has been altered. The Chief Executive of Ziff Messe & Konferenz, Sharyl Leifeld explains: "The exhibition centre in Berlin couldn't allocate us an earlier time slot. The end of June is the middle of the school holidays in many of the German federal states and what is more, these dates conflict with a network trade fair in Great Britain. Our exhibitors prefer an earlier time slot."

In addition to this, Berlin has made the construction times for the stands even shorter for 1995 although the period of time allowed in 1994 was already a bone of contention amongst the exhibitors.

Leifeld goes on to say: "So we've decided to move to Frankfurt am Main. As an important commercial centre, this city lies easy to get to—even by train—and has everything that a modern, professional exhibition site can offer. Starting in 1995, all future German NetWorld+Interop shows will take place in Frankfurt."

Earlier time slot

At the same time, the dates have been put forward. NetWorld+Interop 95 Frankfurt will not take place at the end of June, but rather from the 31st of May to the 2nd of June (the 29th and 30th of May are set aside for tutorials and conferences).

Larger exhibition

The exhibition floor area will be enlarged from 20,000 to 26,000 square meters, of which almost 70 percent are already booked. "This proves the clear acceptance of these changes by the professionals and exhibitors in this sector," says Leifeld. She expects a much larger number of exhibitors, tutorial and conference participants and visitors at NetWorld+Interop 95 in Frankfurt than were present at the premiere in Berlin.

Ziff Messe & Konferenz GmbH, founded in 1993 in Munich, Germany, is a daughter company of Ziff-Davis Exposition and Conference Company (ZD Expos), a conference and trade fair company specialising in computers and communication and based in Foster City, California. ZD Expos organizes several trade shows and conferences worldwide including *NetWorld+Interop, Seybold, Digital World* and *Windows Solutions.*

Letter to the Editor

Ole,

I enjoyed the article on "Routing Arbiter Architecture" in the August 1994 issue of *ConneXions*. Readers might find it useful to know that reference [7], the ISO/IEC specification of IDRP, is available in ASCII text format from any of the standard Internet-Draft repositories as `<draft-kunzinger-idrp-ISO10747-00.txt>`.

—Lyman Chapin
BBN Communications

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use this address for letters to the Editor, questions about back issues etc.:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Suite 201
Foster City, CA 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: `connexions@interop.com`

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 502-493-3217 outside the USA. This is the number for our subscription agency, the Cobb Group. Their fax number is +1 502-491-8050. The mailing address for subscription payments is P.O. Box 35840, Louisville, KY 40232–9496.

NetWorld+Interop World Tour 1995

Las Vegas, NV:	March 27–31
Frankfurt, Germany:	May 29–June 2
Tokyo, Japan:	July 17–21
Atlanta, GA:	September 25–29
Paris, France:	November 6–10

See you there!

conneXions

303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

ADDRESS CORRECTION
REQUESTED

Printed on recycled paper

CONNEXIONS

Subscribe to CONNEXIONS

20% discount prices good until December 31, 1994:

U.S./Canada   ❏ $120 for 1 year (12 issues)   ❏ $210 for 2 years (24)   ❏ $290 for 3 years (36)

International $ 50. additional **per year**   **(Please apply to all of the above.)**

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone (   ) _____

❏ Check enclosed (in U.S. dollars made payable to **CONNEXIONS**).
❏ Visa ❏ MasterCard ❏ American Express ❏ Diners Club   Card # _____ Exp.Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

Back issues available upon request $15./each
Volume discounts available upon request

303 Vintage Park Drive, Suite 201
Foster City, CA 94404-1138
415-578-6900   FAX: 415-525-0194
connexions@interop.com